Seventh Framework Programme – Theme 3.6: Computing Systems

Grant agreement n°248776

D6.4
"Publication Report – Final"

WP6: "Exploitation & Dissemination"

Author: Christian FABRE, CEA.

Date: Saturday 22nd December, 2012

1	CEA (Coord.)	Commissariat à l'énergie atomique et aux énergies alternatives, Laboratoire d'électronique et des technologies de l'information
2	UJF	Université Joseph Fourier Grenoble 1 – VERIMAG
3	ETHZ	Eidgenössische Technische Hochschule Zürich
4	UNIBO	Università di Bologna
5	STM	STMicroelectronics
6	EPFL	École polytechnique Fédérale de Lausanne

Version: v1.0 – Status: Public

CEA ref.: LETI/DACLE/12-0926

Versions of the Document

Version	Date	Author	Comment
0.1	Dec. 9th, 2012	C. FABRE	Initial Version
0.2	Dec. 21st, 2012	I. LLOPARD	Fixed classification of publications
1.0	Dec. 22nd, 2012	C. FABRE	Added the booth at DATE 2012.

Contributors

Christian FABRE (CEA), Ivan LLOPARD (CEA)

Contents

1 Introduction

The objectives of this report is to provide a summary of PRO3D publications across the project's lifetime.

Section 2 below provides some statistics, while each publication, with its abstract, is presented individually in the remaining of this report.

2 Publication Accepted from January 2010 until December 2012

As of December 20th, 2012, PRO3D had more than a hundred and ten publications over the three years of the project.

2.1 Publications per Date

The publication breakdown per year is the following:

- 19 publications in 2010: [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19].
- 56 publication for 2011: [20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75].
- 31 publications for 2012, the last year of the project: [76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100, 103, 104, 105, 106, 107, 108].
- 7 publications, at least, after the end of the project: [109, 110, 112, 113, 114, 115, 116].

2.2 Publications per Kind

The publication breakdown per kind is the following:

- 92 contributions to international conferences and workshops: [1, 2, 3, 4, 5, 6, 7, 8, 9, 11, 12, 13, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 28, 29, 30, 31, 32, 33, 34, 35, 36, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 55, 56, 57, 58, 59, 60, 62, 63, 64, 65, 67, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 87, 88, 89, 91, 93, 94, 95, 96, 99, 100, 103, 104, 105, 106, 107, 108, 114].
- 21 articles in international journals [10, 14, 27, 37, 38, 53, 54, 61, 66, 68, 86, 90, 92, 97, 98, 109, 110, 112, 113, 115, 116].

2.3 Conferences & Workshop per Location

The breakdown of the 92 PRO3D contributions to events (conferences or workshop) is the following:

- 47 contributions to events hosted in Europe [1, 8, 11, 16, 17, 18, 19, 21, 22, 24, 25, 26, 34, 35, 36, 41, 42, 43, 45, 46, 49, 52, 55, 57, 59, 60, 65, 67, 69, 70, 71, 72, 73, 74, 76, 77, 78, 82, 83, 84, 87, 91, 93, 94, 99, 103, 114].
- 1 contributions to events hosted in Middle East or Africa [89].
- 28 contributions to events hosted in North America [2, 4, 5, 6, 7, 9, 12, 13, 15, 30, 32, 33, 44, 50, 51, 56, 58, 62, 63, 64, 79, 81, 85, 88, 95, 96, 100, 107].

- 5 contributions to events hosted in South Americas [23, 28, 29, 106, 108].
- 11 contributions to events hosted in Asia [3, 20, 31, 39, 40, 47, 48, 75, 80, 104, 105].

2.4 Booth at DATE 2012

The project took a booth at DATE 2012 in Dresden – `http://www.date-conference.com/date12` – to present the Y2 results of the project. The booth was open from March 13th to March 15th. Besides flyers and other material, the PRO3D results were illustrated by the poster shown Fig 1 page 6.

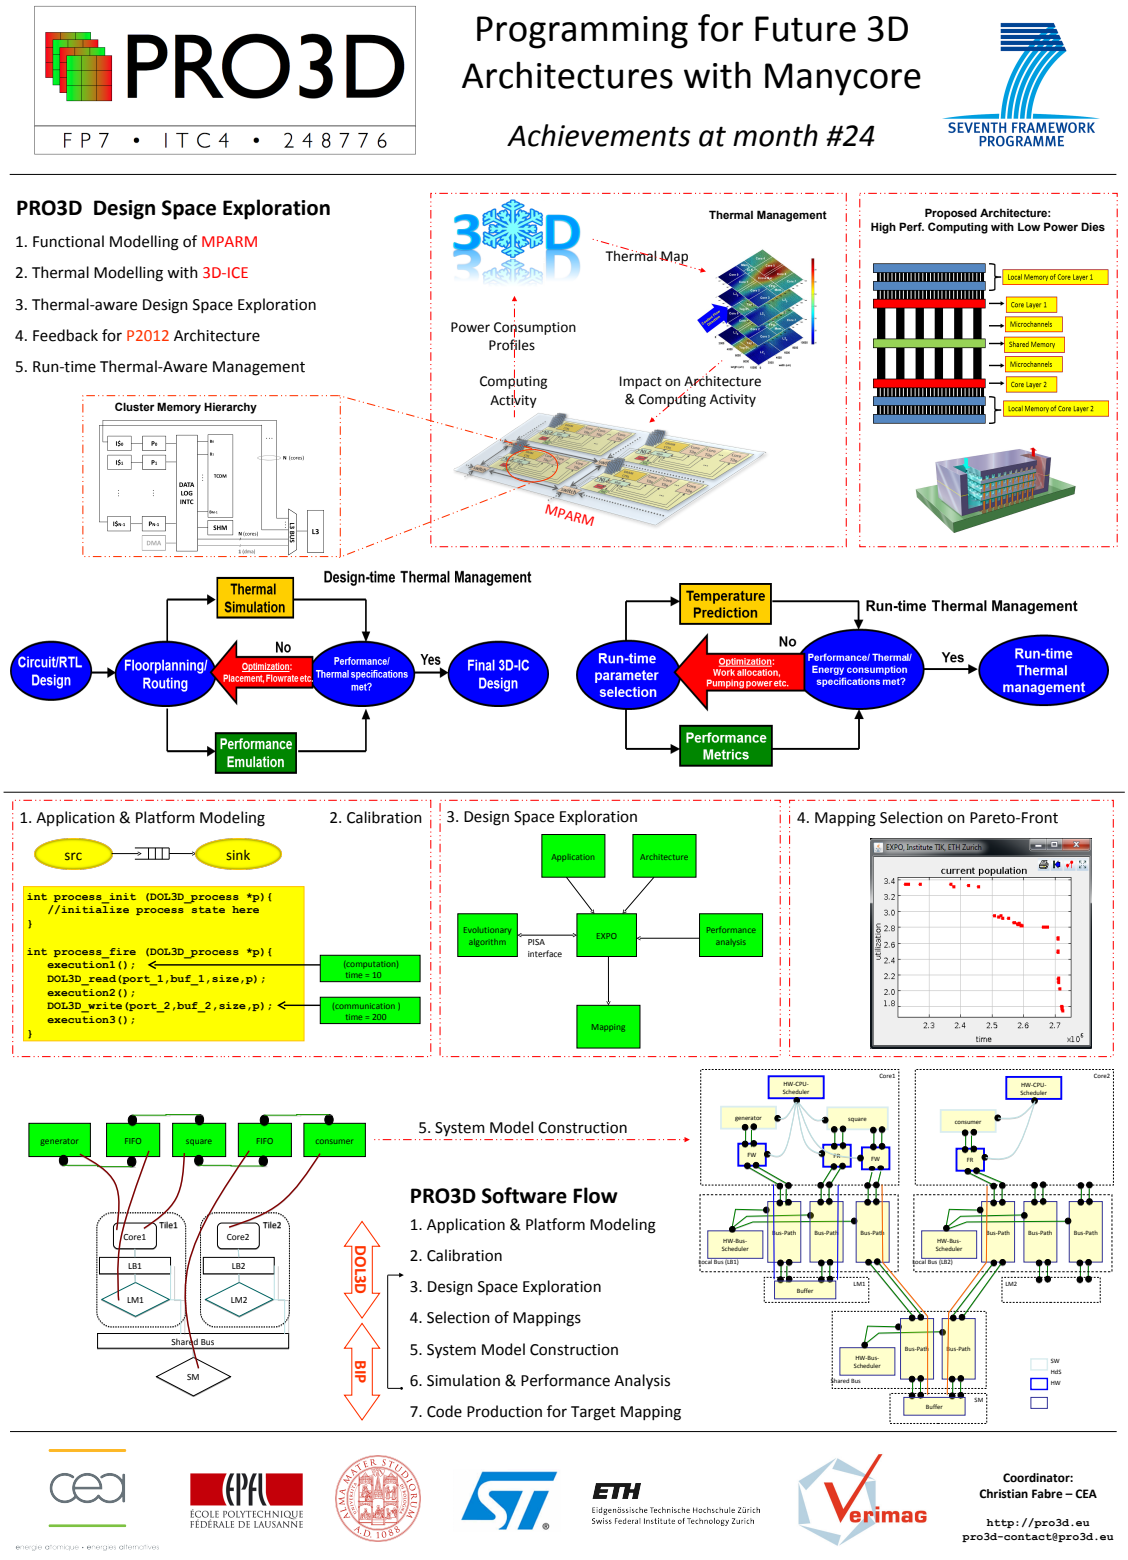

Figure 1: Poster Presented at DATE 2012 – Dresden, Germany – March 12th-16th, 2012

A References

Scientific Contributions Published During the 1st Reporting Period (2010)

[1] *"Automated Conflict-free Distributed Implementation of Component-Based Models"*. B. Bonakdarpour, M. Bozga, M. Jaber, J. Quilbeuf and J. Sifakis In Proceedings of IEEE Fifth International Symposium on Industrial Embedded Systems - SIES 2010, University of Trento, Italy, July 7-9, 2010, pp. 108–117.

We propose a method for generating distributed implementations from high-level models expressed in terms of a set of components glued by rendezvous interactions. The method is a 2-phase transformation preserving all functional properties. The first phase is a source-to-source transformation from global state to a partial state model (to relax atomicity). This transformation replaces multiparty rendezvous interactions by `send()/receive()` *primitives managed by a set of automatically generated distributed schedulers. These schedulers are conflict-free by construction in the sense that they do not require communication in order to safely execute interactions of the high-level model. In the second phase, from the transformed model in phase one, we generate C++ distributed code using either TCP sockets or MPI to implement* `send()/receive()` *primitives. Our method is fully implemented in a tool for automatic generation of distributed applications. We present experimental results using different case studies.*

[2] *"From High-Level Component-Based Models to Distributed Implementations"*. B. Bonakdarpour, M. Bozga, M. Jaber, J.Quilbeuf & J. Sifakis. In Proceedings of the 10th International conference on Embedded software, EMSOFT 2010, Scottsdale, Arizona, USA, October 24-29, 2010, pp. 209–218.

Although distributed systems are widely used nowadays, their implementation and deployment is still a time-consuming, error-prone, and hardly predictive task. In this paper, we propose a methodology for producing automatically efficient and correct-by-construction distributed implementations by starting from a high-level model of the application software in BIP. BIP (Behavior, Interaction, Priority) is a component-based framework with formal semantics that rely on multi-party interactions for synchronizing components and dynamic priorities for scheduling between Interactions. Our methodology transforms arbitrary BIP models into Send/Receive BIP models, directly implementable on distributed execution platforms. The transformation consists of (1) breaking atomicity of actions in atomic components by replacing strong synchronizations with asynchronous `send()/receive()` *interactions; (2) inserting several distributed controllers that coordinate execution of interactions according to a user-defined partition, and (3) augmenting the model with a distributed algorithm for handling conflicts between controllers. The obtained Send/Receive BIP models are proven observationally equivalent to the initial models. Hence, all the functional properties are preserved by construction in the implementation. Moreover, Send/Receive BIP models can be used to automatically derive distributed implementations. Currently, it is possible to generate stand-alone C++ implementations using either TCP sockets for conventional communication, or MPI implementation, for deployment on multicore platforms. This method is fully implemented. We report concrete results obtained under different scenarios (i.e., partitioning of the interactions and choice of algorithm for distributed conflict resolution).*

[3] *"Methods for Knowledge Based Controlling of Distributed Systems"*. S. Bensalem, M. Bozga, S. Graf, D. Peled & S. Quinton. International Symposium on Automated Technol-

ogy for Verification and Analysis, ATVA 2010, September 2010, Singapore, LNCS 6252, pages 52-66.

Controlling concurrent systems to impose some global invariant, is an undecidable problem. One can gain decidability at the expense of reducing concurrency. Even under this flexible design assumption, the synthesis problem remains highly intractable. One practical method for designing controllers is based on checking knowledge properties upon which the processes can make their decisions whether to allow or block transitions. A major deficiency of this synthesis method lies in calculating the knowledge based on the system that we want to control, and not on the resulted system. The original system has less knowledge, and as a result, we may introduce far more synchronization than needed. In this paper we show techniques to reduce this overhead.

[4] *"Integrating Architectural Constraints in Application Software Using Model Transformations in BIP"* Ananda Basu, Saddek Bensalem, Paraskevas Bourgos, Marius Bozga & Joseph Sifakis. Invited presentation at IEEE International High Level Design Validation and Test Workshop. Anaheim, USA. June 10th-12th 2010.

Performance of embedded applications strongly depends on features of the hardware platform on which they are deployed. A grand challenge in the design of complex embedded systems is developing methods and tools for modeling and analyzing the behavior of an application software running on a given hardware architecture. We propose a rigorous method that allows to obtain a model which faithfully represents the behavior of a mixed hardware/software system from a model of its application software and a model of its underlying hardware architecture. The method takes a model of the application software in BIP, a model of the hardware architecture in XML and a mapping associating read and write operations of the application software with execution paths in the architecture. It builds a model of the corresponding system in BIP. The latter can be simulated and analyzed for the verication of both functional and extra-functional properties. The method consists in progressively enriching the application software model. It involves two steps: (1) The identication of the a set of hardware components used in the system model. These are included in the initial model. They are obtained by a systematic decomposition of the architecture through an abstract grammar. (2). The application of a sequence of source-to-source transformations to synthesize run-time routines implementing software channels. The transformations are correct-by-construction. In particular they preserve functional properties of the application software. We have identied a minimal and complete set of transformation rules, that are proved to be correct. The system model is highly parametrized and allows exible integration of specic target architecture features, such as bus policy and scheduling policy of the processor cores The method has been implemented for software applications and hardware architectures described in the DOL tool for performance evaluation. We plan to test benchmark applications on a simplied MPARM architecture.

[5] *"Thermal-Aware Compilation for System-on-Chip Processing Architectures"*. Aly, Mohamed M. Sabry ; Rodrigo, Ayala ; Luis, José ; Atienza Alonso, David. Proceedings of the 20th ACM Great Lakes Symposium on VLSI (GLSVLSI 2010), vol. 1, num. 1, 2010, p. 221–226, New York: ACM Press. ISBN: 978-1-4503-0012-4/10/06. Providence, Rhode Island, USA. May 16-18 2010.

The development of compiler-based mechanisms to reduce the percentage of hotspots and optimize the thermal profile of large register files has become an important issue. Thermal hotspots have been known to cause severe reliability issues, while the thermal profile of the devices is also related to the leakage power consumption and the cooling cost. In this paper we propose several compilation techniques that, based on an efficient register allocation mechanism, reduce the percentage of hotspots in the register file and uniformly distribute the heat. As a result, the thermal profile and

reliability of the device is clearly improved. Simulation results show that the proposed flow achieved 91% reduction of hotspots and 11% reduction of the peak temperature.

[6] *"Performance and Energy Trade-offs Analysis of L2 on-Chip Cache Architectures for Embedded MPSoCs"*. Aly, Mohamed M. Sabry ; Ruggiero, Martino ; Garcia del Valle, Pablo. Proceedings of the 20th symposium on Great lakes symposium on VLSI, 2010, p. 305–310, ISBN: 978-1-4503-0012-4. Providence, Rhode Island, USA. May 16th-18th 2010.

On-chip memory organization is one of the most important aspects that can influence the overall system behavior in multi-processor systems. Following the trend set by high-performance processors, high-end embedded cores are moving from single-level on chip caches to a two-level on-chip cache hierarchy. Whereas in the embedded world there is general consensus on L1 private caches, for L2 there is still not a dominant architectural paradigm. Cache architectures that work for high performance computers turn out to be inefficient for embedded systems (mainly due to power-efficiency issues). This paper presents a virtual platform for design space exploration of L2 cache architectures in low-power Multi-Processor-Systems-on-Chip (MPSoCs). The tool contains several L2 caches templates, and new architectures can be easily added using our flexible plugin system. Given a set of constrains for a specific system (power, area, performance), our tool will perform extensive exploration to find the cache organization that best suits our needs. Through some practical experiments, we show how it is possible to select the optimal L2 cache, and how this kind of tool can help designers avoid some common misconceptions. Benchmarking results in the experiments section will show that for a case study with multiple processors running communicating tasks allocated on different cores, the private L2 cache organization still performs better than the shared one.

[7] *"3D-ICE: Fast Compact Transient Thermal Modeling for 3D-ICs with Inter-Tier Liquid Cooling"*. Sridhar, Mahankali ; Raj, Arvind ; Vincenzi, Alessandro ; Ruggiero, Martino ; Brunschwiler, Thomas ; Atienza Alonso, David. Proceedings of the 2010 International Conference on Computer-Aided Design (ICCAD 2010), vol. 1, num. 1, 2010, p. 1-8. New York: ACM and IEEE Press, 2010. San Jose, California, USA. November 7th-11th, 2010.

Three dimensional stacked integrated circuits (3D ICs) are extremely attractive for overcoming the barriers in interconnect scaling, offering an opportunity to continue the CMOS performance trends for the next decade. However, from a thermal perspective, vertical integration of high-performance ICs in the form of 3D stacks is highly demanding since the effective areal heat dissipation increases with number of dies (with hotspot heat fluxes up to 250W/cm^2) generating high chip temperatures. In this context, inter-tier integrated microchannel cooling is a promising and scalable solution for high heat flux removal. A robust design of a 3D IC and its subsequent thermal management depend heavily upon accurate modeling of the effects of liquid cooling on the thermal behavior of the IC during the early stages of design. In this paper we present 3D-ICE, a compact transient thermal model (CTTM) for the thermal simulation of 3D ICs with multiple inter-tier microchannel liquid cooling. The proposed model is compatible with existing thermal CAD tools for ICs, and offers significant speed-up (up to × 975) over a typical commercial computational fluid dynamics simulation tool while preserving accuracy (i.e., maximum temperature error of 3.4%). In addition, a thermal simulator has been built based on 3D-ICE, which is capable of running in parallel on multicore architectures, offering further savings in simulation time and demonstrating efficient parallelization of the proposed approach.

[8] *"Compact Transient Thermal Model for 3D ICs with Liquid Cooling via Enhanced Heat Transfer Cavity Geometries"*. Sridhar, Mahankali ; Raj, Arvind ; Vincenzi, Alessandro ; Ruggiero, Martino ; Brunschwiler, Thomas ; Atienza Alonso, David. Proceedings of

the 16th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC'10), vol. 1, num. 1, 2010, p. pp. 5 – 11. New York: IEEE Press, 2010. Barcelona, Spain. October 6th-8th, 2010.

The advent of 3D stacked ICs with accumulating heat fluxes stresses thermal reliability and is responsible for temperature driven performance deterioration of the electronic systems Hot spots with power densities typically rising up to 250 W/cm^2 are not acceptable, with the result of limited performance improvement in next generation high performance microprocessor stacks. Unfortunately traditional back-side cooling only scales with the chip stack footprint, but not with the number of tiers. Direct heat removal from the IC dies via inter-tier liquid cooling is a promising solution to address this problem. In this regard, a thermal-aware design of a 3D IC with liquid cooling for optimal electronic performance and reliability requires fast modeling and simulation of the liquid cooling during the early stages of the design. In this paper, we propose a novel compact transient thermal modeling (CTTM) scheme for liquid cooling in 3D ICs via microchannels and enhanced heat transfer cavity geometries such as pin-fin structures. The model is compatible with the existing thermal-CAD tools for ICs and offers significant speed-up over commercial computational fluid dynamics simulators (× 13478 for pin-fin geometry with 1.1% error in temperature). In addition, the model is highly flexible and it provides a generic framework in which heat transfer coefficient data from numerical simulations or existing correlations can be incorporated depending upon the speed/accuracy needs of the designer. We have also studied the effects of using different techniques for the estimation of heat transfer coefficients on the accuracy of the model. This study highlights the need to consider developing flow conditions to accurately model the temperature field in the chip stack. The use of correlation data from fully developed flows only results in temperature error as high as 9 K (about 41%) near the inlet.

[9] *"Fuzzy Control for Enforcing Energy Efficiency in High-Performance 3D Systems"*. Sabry Aly, Mohamed Mostafa, Ayse Kivilcim Coskun, David Atienza Alonso. Proceedings of the 2010 International Conference on Computer-Aided Design (ICCAD 2010), vol. 1, num. 1, 2010, p. 10-16. New York: IEEE Press, 2010. San Jose, California, USA. November 7th-11th, 2010.

3D stacked circuits reduce communication delay in multicore system-on-chips (SoCs) and enable heterogeneous integration of cores, memories, sensors, and RF devices. However, vertical integration of layers exacerbates the reliability and thermal problems, and cooling is a limiting factor in multitier systems. Liquid cooling is a highly efficient solution to overcome the accelerated thermal problems in 3D architectures; however, liquid cooling brings new challenges in modeling and runtime management. This paper proposes a novel controller for improving energy efficiency and reliability in 3D systems through liquid cooling management and dynamic voltage frequency scaling (DVFS). The proposed fuzzy controller adjusts the liquid flow rate at runtime to match the cooling demand for preventing energy wastage of over-cooling and for maintaining a stable thermal profile. The DVFS decisions provide chip-level energy savings and help balancing the temperature across the system. Experimental results on 8- and 16-core multicore SoCs show that the controller prevents the system to exceed the given threshold temperature while reducing cooling energy by up to 50% and system-level energy by up to 21% in comparison to using a static worst-case flow rate setting.

[10] *"Thermal-Aware Compilation for Register Window-Based Embedded Processors"*. Aly, Sabry ; Mostafa, Mohamed ; Rodrigo, Ayala ; Luis, José ; Atienza Alonso, David. IEEE Embedded Systems Letters, vol. 2, num. 4, 2010, p. 213-217, Institute of Electrical and Electronics Engineers, 2010. ISSN: 1943-0663.

The development of compiler-based mechanisms to optimize the thermal profile of large register files to improve the processor reliability has become an important issue. Thermal hotspots have been known to cause severe reliability issues, while the thermal profile of the devices is also related to the leakage power consumption and the cooling cost. Register window-based architectures provide a relatively large register files. However, such large register files are not designed or utilized for thermal balancing or reliability enhancement. In this paper, we propose a compilation flow that utilizes the register windows to reduce optimize the thermal profile and to reduce the hotspots. As a result, the thermal profile and reliability of the device is clearly improved. Simulation results show that the proposed flow achieves up to 91% reduction of hotspots and 11% reduction of the peak temperature in embedded processors.

[11] *"Evaluating OpenMP Support Costs on MPSoCs"*. Marongiu A., Burgio P., Benini L.; Proceedings of the 13th Euromicro Conference on Digital Systems Design (DSD 2010). Lille, France. September 1st-3rd, 2010.

The ever-increasing complexity of MPSoCs is making the production of software the critical path in embedded system development. Several programming models and tools have been proposed in the recent past that aim at facilitating application development for embedded MPSoCs. OpenMP is a mature and easy-to-use standard for shared memory programming, which has recently been successfully adopted in embedded MPSoC programming as well. To achieve performance, however, it is necessary that the implementation of OpenMP constructs efficiently exploits the many peculiarities of MPSoC hardware. In this paper we present an extensive evaluation of the cost associated with supporting OpenMP on such a machine, investigating several implementative variants that efficiently exploit the memory hierarchy. Experimental results on different benchmarks confirm the effectiveness of the optimizations in terms of performance improvements.

[12] *"Vertical Stealing: Robust, Locality-Aware Do-All Workload Distribution for 3D MPSoCs"*; A. Marongiu, P. Burgio, L. Benini. Proceedings of the 2010 International Conference on Compilers, Architecture, and Synthesis for Embedded Systems. October 24th-28th, 2010. Scottsdale, Arizona, USA.

In this paper we address the issue of efficient doall workload distribution on a embedded 3D MPSoC. 3D stacking technology enables low latency and high bandwidth access to multiple, large memory banks in close spatial proximity. In our implementation one silicon layer contains multiple processors, whereas one or more DRAM layers on top host a NUMA memory subsystem. To obtain high locality and balanced workload we consider a two-step approach. First, a compiler pass analyzes memory references in a loop and schedules each iteration to the processor owning the most frequently accessed data. Second, if locality-aware loop parallelization has generated unbalanced workload we allow idle processors to execute part of the remaining work from neighbors by implementing runtime support for work stealing.

[13] *"Exploring Programming Model-Driven QoS Support for NoC-based Platforms"*. Joven J., Marongiu A., Angiolini F., Benini L., De Micheli G.; Proceedings of the 2010 International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS 2010). Scottsdale, Arizona, USA. October 24th-29th, 2010.

Networks-on-Chip (NoCs) are being increasingly considered as a central enabling technology to communication-centric designs as more and more IP blocks are integrated on the same SoC. Embedded applications, in turn, are becoming extremely sophisticated, and often require guaranteed levels of service and performance. The complex and non-uniform nature of network traffic generated by parallel applications running on a large number of possibly heterogeneous IPs makes a strong case for providing Quality of Service (QoS) support for traffic streams over the NoC infrastructure. In this paper we consider an integrated hardware/software approach for delivering

QoS at the application level. We designed NoC hardware support, low-level middleware and APIs which enable QoS control at the application level. Furthermore, we identify a set of programming abstractions useful to associate the notion of priority to each running task in the system. An initial implementation of this programming model is also presented, which leverages a set of extensions to a MPSoC-specific OpenMP compiler and runtime environment.

[14] *"An OpenMP Compiler for Efficient Use of Distributed Scratchpad Memory in MPSoCs"*. Marongiu A., Benini L.; IEEE Transactions on Computers, Vol. 59, Issue 12. December 2010. DOI: 10.1109/TC.2010.199.

Most of today's state-of-the-art processors for mobile and embedded systems feature on-chip scratchpad memories. To efficiently exploit the advantages of low-latency high-bandwidth memory modules in the hierarchy there is the need for programming models and/or language features that expose such architectural details. On the other hand, effectively exploiting the limited on-chip memory space requires the programmer to devise an efficient partitioning and distributed placement of shared data at the application level. In this paper we propose a programming framework that combines the ease of use of OpenMP with simple yet powerful language extensions to trigger array data partitioning. Our compiler exploits profiled information on array access count to automatically generate data allocation schemes optimized for locality of references.

[15] *"Performance and Energy Trade-offs Analysis of L2 on-Chip Cache Architectures for Embedded MPSoCs"*. M. M. S. Aly, M. Ruggiero and P. Garcia del Valle. Proceedings of the 20th symposium on Great lakes symposium on VLSI, Providence, Rhode Island, USA. May 16th-18th, 2010

On-chip memory organization is one of the most important aspects that can in°uence the overall system behavior in multi-processor systems. Following the trend set by high-performance processors, high-end embedded cores are moving from single-level on chip caches to a two-level on-chip cache hierarchy. Whereas in the embedded world there is general consensus on L1 private caches, for L2 there is still not a dominant architectural paradigm. Cache architectures that work for high performance computers turn out to be ine±cient for embedded systems (mainly due to power-e±ciency issues). This paper presents a virtual platform for design space exploration of L2 cache architectures in low-power Multi-Processor-Systems-on-Chip (MPSoCs). The tool contains several L2 caches templates, and new architectures can be easily added using our flexible plug-in system. Given a set of constrains for a specific system (power, area, performance), our tool will perform extensive exploration to find the cache organization that best suits our needs. Through some practical experiments, we show how it is possible to select the optimal L2 cache, and how this kind of tool can help designers avoid some common misconceptions. Benchmarking results in the experiments section will show that for a case study with multiple processors running communicating tasks allocated on different cores, the private L2 cache organization still performs better than the shared one.

[16] *"Scalable Instruction Set Simulator for Thousand-Core Architectures Running on GPGPUs"*. Shivani Raghav, Martino Ruggiero, David Atienza, Christian Pinto, Andrea Marongiu, Luca Benini. Proc. of the 2010 International Conference on High Performance Computing & Simulation (HPCS 2010), pages 459–466. June 28th-July 2nd, 2010. Caen, France.

Simulators are still the primary tools for development and performance evaluation of applications running on massively parallel architectures. However, current virtual platforms are not able to tackle the complexity issues introduced by 1000-core future scenarios. We present a fast and accurate simulation framework targeting extremely large parallel systems by specifically taking advantage of the inherent potential processing parallelism available in modern GPGPUs.

[17] *"Programming Model and Compiler Support for Efficient Data Distribution in PGAS MPSoCs"*. Andrea Marongiu & Luca Benini. Poster presented at Intel European Research and Innovation Conference (ERIC). September 21th-22nd, 2010. Braunschweig, Germany. `http://www.intel.com/corporate/education/emea/event/irc/`. See Fig. 2 page 44.

[18] *"Efficient Execution of Kahn Process Networks on CELL BE"*. Iuliana Bacivarov, Wolfgang Haid, Kai Huang, Lars Schor & Lothar Thiele. Presented at the Summer School on Models for Embedded Signal Processing Systems. Lorentz Center, Leiden, The Netherlands. August 30th-September 3rd, 2010. `http://www.lorentzcenter.nl/lc/web/2010/427/presentations/Iuliana-cell.pdf`.

[19] *"Distributed Operation Layer... A Practical Perspective"*. Iuliana Bacivarov, Wolfgang Haid, Kai Huang & Lothar Thiele. Presented at the Summer School on Models for Embedded Signal Processing Systems. Lorentz Center, Leiden, The Netherlands. August 30th-September 3rd, 2010. `http://www.lorentzcenter.nl/lc/web/2010/427/presentations/Iuliana-demo.pdf`.

Scientific Contributions Published During the 2nd Reporting Period (2011)

[20] *"Thermally Optimal Stop-Go Scheduling of Task Graphs with Real-Time Constraints"*. Pratyush Kumar, Lothar Thiele. Asia and South Pacific Design Automation Conference. Jan. 25th-28th 2011, Yokohama, Japan.

Dynamic thermal management (DTM) techniques to manage the load on a system to avoid thermal hazards are soon becoming mainstream in today's systems. Several research works have studied the use of dynamic voltage scaling (DVS) to serve a set of independent real-time tasks with temperature/performance constraints. In this paper, we study another important class of DTM techniques: stop-go scheduling, to minimize peak temperature when scheduling an application modeled as a task-graph within a given makespan constraint. For a given static-ordering of execution of tasks, we derive the optimal schedule referred to as a JUST policy. We also prove that for periodic task-graphs the optimal temperature is independent of the chosen static-ordering. Simulation experiments validate the theoretical results.

[21] *"Worst-Case Temperature Analysis for Real-Time Systems"*. Lothar Thiele, Iuliana Bacivarov, Jian-Jia Chen, Devendra Rai, Hoeseok Yang. Proc. of Design, Automation and Test in Europe (DATE '11), ACM and IEEE Press. Grenoble, France. March 14th-18th, 2011.

With the evolution of today's semiconductor technology, chip temperature increases rapidly mainly due to the growth in power density. For modern embedded real-time systems, it is crucial to estimate maximal temperatures in order to take mapping or other design decisions to avoid burnout, and still be able to guarantee meeting real-time constraints. This paper provides answers to the question: When work-conserving scheduling algorithms, such as earliest-deadline first (EDF), rate-monotonic (RM), deadline-monotonic (DM), are applied, what is the worst-case peak temperature of a real-time embedded system under all possible scenarios of task executions? We propose an analytic framework, which considers a general event model based on network and real-time calculus. This analysis framework has the capability to handle a broad range of uncertainties in terms of task execution times, task invocation periods, and jitter in task arrivals. Simulations show that our framework is a cornerstone to design real-time systems that have guarantees on both schedulability and maximal temperatures.

[22] *"Towards Thermally-Aware Design of 3D MPSoCs with Inter-Tier Cooling"*, Mohamed M. Sabry, Arvind Sridhar, David Atienza, Yuksel Temiz, Yusuf LeblebicI, Sylwia Szczukiewicz, Navid Borhani, John R. Thome, Thomas Brunschwiler, and Bruno Michel. Proc. of Design, Automation and Test in Europe (DATE '11), ACM and IEEE Press. Grenoble, France. March 14th-18th, 2011.

New tendencies envisage 3D Multi-Processor System-On-Chip (MPSoC) design as a promising solution to keep increasing the performance of the next-generation highperformance computing (HPC) systems. However, as the power density of HPC systems increases with the arrival of 3D MPSoCs, supplying electrical power to the computing equipment and constantly removing the generated heat is rapidly becoming the dominant cost in any HPC facility. Thus, both power and thermal/cooling implications play a major role in the design of new HPC systems, given the energy constraints in our society. Therefore, in this work we propose a new holistic thermally-aware design. This paper presents the exploration of novel cooling technologies, as well as suitable thermal modeling and system-level design methods, which are all necessary to develop 3D MPSoCs with inter-tier liquid cooling systems. As a result, we develop energy-efficient run-time thermal control strategies to achieve energy-efficient cooling mechanisms to compress almost 1 Tera nano-sized functional units into one cubic centimeter with a 10 to 100 fold higher connectivity than otherwise possible. The proposed thermally-aware design paradigm includes exploring the synergies of hardware-, software- and mechanical-based thermal control techniques as a fundamental step to design 3D MPSoCs for HPC systems. Our management strategy prevents the system from surpassing the given threshold temperature while achieving up to 67% reduction in cooling energy and up to 30% reduction in system-level energy in comparison to setting the flow rate at the maximum value to handle the worst-case temperature.

[23] *"3D Thermal-Aware Floorplanner for Many-Core Single-Chip Systems"*, David Cuesta, Jose L. Risco, Jose L. Ayala, David Atienza, Proc. of 12th IEEE Latin-American Test Workshop (LATW11), IEEE Press. Beach of Porto de Galinhas, Brazil. March 27th-30th, 2011.

Heat removal and power density distribution delivery have become two major reliability concerns in 3D stacked technology. In this paper, we propose a thermal-driven 3D floorplanner. Our contributions include: (1) a novel multi-objective formulation to consider the thermal and performance constraints in the optimization approach; (2) an efficient Mixed Integer Linear Programming (MILP) representation of the floorplanning model; and (3) a smooth integration of the MILP model with an accurate thermal modelling of the architecture. The experimental work is conducted for two realistic many-core single-chip architectures: an homogeneous system resembling Intel's SCC, and an improved heterogeneous setup. The results show promising improvements of the mean, peak temperature and the thermal gradient, with a reduced overhead in the wire length of the system.

[24] *"Thermal-Aware System Analysis and Software Synthesis for Embedded Multi-Processors"*, Lothar Thiele, Lars Schor, Hoeseok Yang and Iuliana Bacivarov. Design Automation & Test in Europe. March 14th-18th, 2011. Grenoble, France. `ftp://ftp.tik.ee.ethz.ch/pub/people/lschor/tsyb2011a.pdf`.

Nowadays, the reliability and performance of modern embedded multi-processor systems is threaten by the ever-increasing power densities in integrated circuits, and a new additional goal of software synthesis is to reduce the peak temperature of the system. However, in order to perform thermal-aware mapping optimization, the timing and thermal characteristics of every candidate mapping have to be analyzed. While the task of analyzing timing characteristics of design alternatives has been extensively investigated in recent years, there is still a lack of methods for accurate and fast thermal analysis. In order to obtain desired evaluation times, the system has to be simulated at a

high abstraction level. This often results in a loss of accuracy, mainly due to missing knowledge of system's characteristics. This paper addresses this challenge and presents methods to automatically calibrate high-level thermal evaluation methods. Furthermore, the viability of the methods for automated model calibration is illustrated by means of a novel high-level thermal evaluation method.

[25] *"Distributed Application Layer: Towards Efficient and Reliable Programming of Many-Tile Architectures".* Iuliana Bacivarov, Hoeseok Yang, Lars Schor, Devendra Rai, Sudhanshu Jha, and Lothar Thiele. Design Automation & Test in Europe, March 14th-18th, 2011. Grenoble, France. `http://conferenze.dei.polimi.it/depcp/proceedings/pdfs/3.pdf`.

Many-tile architectures are large lattices that have multi-processor system-on-chips (MPSoC) as components, and as such, they push the MPSoC concept one step further in terms of raw computational power and homogeneous/ heterogeneous application parallelism. This advanced degree of computational flexibility comes at expenses of new requirements in terms of programming paradigms. The distributed operation layer (DOL) (http://www.tik.ee.ethz.ch/ shapes/) was successfully utilized for programming and optimizing the static mapping of one parallel application on different multi-processor platforms like PC clusters, MPARM, IBM Cell, and Atmel D940. However, to cope with multiple, dynamic, and concurrent applications and to optimize their execution while still guaranteeing real-time constraints, a new programming model, named distributed application layer (DAL) is being developed on top of DOL. Besides offering an efficient programming model for many-tile architectures, DAL provides faster application development and reliability awareness. This poster introduces all basic concepts of DAL, presenting the programming model, design principles, and a proof-of-concept implementation.

[26] *"Towards Thermally-Aware Design of 3D MPSoCs with Inter-Tier Cooling"*. Mohamed Sabry, Arvind Sridhar, David Atienza Alonso and Yuksel Temiz. Design Automation & Test in Europe, March 14th-18th, 2011. Grenoble, France.

New tendencies envisage 3D Multi-Processor System-On-Chip (MPSoC) design as a promising solution to keep increasing the performance of the next-generation highperformance computing (HPC) systems. However, as the power density of HPC systems increases with the arrival of 3D MPSoCs, supplying electrical power to the computing equipment and constantly removing the generated heat is rapidly becoming the dominant cost in any HPC facility. Thus, both power and thermal/cooling implications play a major role in the design of new HPC systems, given the energy constraints in our society. Therefore, EPFL, IBM and ETHZ have been working within the CMOSAIC Nano-Tera.ch program project in the last three years on the development of a holistic thermally-aware design. This paper presents the exploration in CMOSAIC of novel cooling technologies, as well as suitable thermal modeling and system-level design methods, which are all necessary to develop 3D MPSoCs with inter-tier liquid cooling systems. As a result, we develop energy-efficient run-time thermal control strategies to achieve energy-efficient cooling mechanisms to compress almost 1 Tera nano sized functional units into one cubic centimeter with a 10 to 100 fold higher connectivity than otherwise possible. The proposed thermally-aware design paradigm includes exploring the synergies of hardware-, software- and mechanical-based thermal control techniques as a fundamental step to design 3D MPSoCs for HPC systems. More precisely, we target the use of inter-tier coolants ranging from liquid water and twophase refrigerants to novel engineered environmentally friendly nano-fluids, as well as using specifically designed micro-channel arrangements, in combination with the use of dynamic thermal management at system-level to tune the flow rate of the coolant in each micro-channel to achieve thermally-balanced 3D-ICs. Our management strategy prevents the system from surpassing the given threshold temperature while achieving up to 67% reduction

in cooling energy and up to 30% reduction in system-level energy in comparison to setting the flow rate at the maximum value to handle the worst-case temperature.

[27] *"Rigorous Component-Based System Design Using the BIP Framework"*. Ananda Basu, Saddek Bensalem, Marius Bozga, Jacques Combaz, Mohamad Jaber, Thanh-Hung Nguyen, Joseph Sifakis. IEEE Software, vol. 28, no. 3, May/June 2011. DOI: `http://dx.doi.org/10.1109/MS.2011.27`.

Rigorous system design requires the use of a single powerful component framework allowing the representation of the designed system at di fferent levels of detail, from application software to its implementation. The use of a single framework allows to maintain the overall coherency and correctness by comparing di fferent architectural solutions and their properties. In this paper, we present the BIP (Behavior, Interaction, Priority) component framework which encompasses an expressive notion of composition for heterogeneous components by combining interactions and priorities. This allows description at di fferent levels of abstraction from application software to mixed hardware/software systems. Then, we introduce a rigorous design flow that uses BIP as a unifying semantic model to derive from an application software, a model of the target architecture and a mapping, a correct implementation. Correctness of implementation is ensured by application of source-to-source transformations in BIP which preserve correctness of essential design properties. The design is fully automated and supported by a toolset including a compiler, the D-Finder veri cation tool and model transformers. We illustrate the use of BIP as a modeling formalism as well as crucial aspects of the design flow for ensuring correctness, through an autonomous robot case study.

[28] *" Thermal-Aware System-Level Modeling and Management for Multi-Processor Systems-on-Chip"*. Francesco Zanini, David Atienza, Luca Benini and Giovanni de Micheli. Proceedings of IEEE International Symposium on Circuits and Systems (ISCAS'11), pages 2481–2484. Rio de Janeiro, Brazil, May 15th-18th, 2011. ISBN: 978-1-4244-9472-9/11. `http://infoscience.epfl.ch/record/162539`.

Multi-Processor Systems-on-Chip (MPSoCs) are penetrating the electronics market as a powerful, yet commercially viable, solution to answer the strong and steadily growing demand for scalable and high performance systems, at limited design complexity. However, it is critical to develop dedicated system-level design methodologies for multi-core architectures that seamlessly address their thermal modeling, analysis and management. In this work, we first formulate the problem of system-level thermal modeling and link it to produce a global thermal management formulation as a discrete-time optimal control problem, which can be solved using finite-horizon model-predictive control (MPC) techniques, while adapting to the actual time-varying unbalanced MPSoC workload requirements. Finally, we compare the system-level MPC-based thermal modeling and management approaches on an industrial 8-core MPSoC design and show their different trade-offs regarding performance while respecting operating temperature bounds.

[29] *"Thermal Analysis and Active Cooling Management for 3D MPSoCs"*. M. Sabry, D. Atienza and A. K. Coskun. Proceedings of IEEE International Symposium on Circuits and Systems (ISCAS'11), pages 2237–2240. Rio de Janeiro, Brazil. May 15th-18th, 2011. ISBN: 978-1-4244-9472-9/11. `http://infoscience.epfl.ch/record/163372`.

3D stacked architectures reduce communication delay in multiprocessor system-on-chips (MPSoCs) and allowing more functionality per unit area. However, vertical integration of layers exacerbates the reliability and thermal problems, and cooling is a limiting factor in multi-tier systems. Liquid cooling is a highly efficient solution to overcome the accelerated thermal problems in 3D architectures. However, liquid cooling brings new challenges in modeling and run-time management. This

paper proposes a design-time/run-time thermal management policy for 3D MPSoCs with inter-tier liquid cooling. First, we perform a design-time analysis to estimate the thermal impact of liquid cooling and dynamic voltage frequency scaling (DVFS) on 3D MPSoCs. Based on this analysis, we define a set of management rules for run-time thermal management. We utilize these rules to control and adjust the liquid flow rate in order to match the cooling demand for preventing energy wastage of overcooling, while maintaining a stable thermal profile in the 3D MPSoCs. Experimental results on multi-tier 3D MPSoCs show that proposed design-time/run-time management policy prevents the system to exceed the given threshold temperature while reducing cooling energy by 50% on average and system-level energy by 18% on average in comparison to using a static worstcase flow rate setting.

[30] *"GPGPU-Accelerated Parallel and Fast Simulation of Thousand-core Platforms"*, Christian Pinto, Shivani Raghav, Andrea Marongiu, Martino Ruggiero, David Atienza, Luca Benini. IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing 2011 (CCGRID2011). Newport Beach, California. May 23rd-26th 2011. DOI: 10.1109/CCGrid.2011.64. `http://infoscience.epfl.ch/record/164471/files/ccgrid11.pdf`.

The multicore revolution and the ever-increasing complexity of computing systems is dramatically changing sys- tem design, analysis and programming of computing platforms. Future architectures will feature hundreds to thousands of simple processors and on-chip memories connected through a Network-on-Chip (NoC). Architectural simulators will remain primary tools for design space exploration, software development and performance evaluation of these massively parallel architectures. However, architectural simulation performance is a serious concern, as virtual platforms and simulation technology are not able to tackle the complexity of thousands of core future scenarios. The main contribution of this paper is the development of a new simulation approach and technology for many core processors which exploit the enormous parallel processing capability of low-cost and widely available General Purpose Graphic Processing Units (GPGPU). The simulation of many-core architectures exhibits indeed a high level of parallelism and is inherently parallelizable, but GPGPU acceleration of architectural simulation requires an in-depth revision of the data structures and functional partitioning traditionally used in parallel simulation. We demonstrate our GPGPU simulator on a target architecture composed by several cores (i.e. ARM ISA based), with instruction and data caches, connected through a NoC. Our experiments confirm the feasibility of our approach.

[31] *"Die-Level TSV Fabrication Platform for CMOS-MEMS Integration"*. Yuksel Temiz, Michail Zervas, Carlotta Guiducci, and Yusuf Leblebici. Proceedings of the 16th International Conference on Solid-State Sensors, Actuators and Microsystems (Transducers 2011). June 5th-9th, 2011. Beijing, China.

This paper reports a new post-CMOS processing platform for die-level through-silicon-via (TSV) fabrication, based on wafer reconstitution from embedded dies, parylene deposition, stencil lithography, and bottom-up electroplating. The goal of this work is to develop a heterogeneous 3D-integration technique for the applications requiring CMOS-MEMS integration with vertical interconnections.

[32] *"Cool shapers: Shaping Real-Time Tasks for Improved Thermal Guarantees"*, Pratyush Kumar and Lothar Thiele, Design Automation Conference 2011, June 5th-10th, 2011. San Diego, California, USA. `ftp://blackbox.ethz.ch/pub/people/kumarpr/DAC11.pdf`

With increasing power densities, managing on-chip temperatures has become an important design challenge. We propose a novel approach to this problem with the use of shapers to dynamically and selectively insert idle times during the execution of hard real-time jobs on a single speed processor.

For the class of leaky bucket shapers which have a light-weight implementation, we derive the shaper such that no job misses its real-time deadline and the peak temperature is optimally reduced. The analysis and design of such shapers allows for dynamically variable streams of jobs; for instance, periodic streams with jitter. We extend our results to consider non-zero power and timing overhead in transitioning to the idle mode. With experimental results, we demonstrate that the proposed approach provides a large improvement: on average 8K peak temperature reduction or 40% increase in utilization for a given peak temperature.

[33] *"Temperature Predictability in Multi-Core Real-Time Systems".* Iuliana Bacivarov, Lars Schor, Hoeseok Yang, Devendra Rai, and Lothar Thiele. Presentation at the DAC Workshop on Multiprocessor System on Chip for Cyber Physical Systems: Programmability, Run-Time Support and Hardware Platforms for High Performance Embedded Applications. At DAC 2011, San Diego, California, USA. June 9th, 2011. `http://www.dac.com/workshops+_+colocated+events.aspx?event=104&topic=15`.

[34] *"Thermal-Aware Design of Real-Time Multi-Core Embedded Systems".* Iuliana Bacivarov, Lars Schor, Hoeseok Yang, Devendra Rai, and Lothar Thiele. Presentation at the Artemis Workshop on Mapping of Applications to MPSoCs (Map2MPSoC 2011). Rheinfels Castle, St. Goar, Germany. June 28th-29th, 2011. `http://www.artist-embedded.org/artist/Program,2298.html`.

[35] *"End-to-End Delay Minimization in Thermally Constrained Distributed Systems"*, Pratyush Kumar and Lothar Thiele, Euromicro Conference on Real-Time Systems (ECRTS 2011). July 6th–8th, 2011. Porto, Portugal.

With ever-increasing power densities, managing on-chip temperatures by optimizing mapping and scheduling of tasks is becoming increasingly necessary. We study the minimization of end-to-end delay for thermally constrained scheduling of an application that is specified as a task graph and is executing on parallel processors without speed scaling. We show that task graph scheduling on thermally constrained systems is monotonic, i.e., delaying the execution of a task longer than necessary cannot lead to the early completion of any other task. Using this monotonicity principle, we design the provably optimal schedule for a given mapping, called the JUST schedule. The JUST schedule can be easily implemented using temperature sensors. We then present different thermal-aware modifications to standard mapping heuristics and evaluate them on a large set of problem instances. The experimental results illustrate that with simple thermal-aware modifications, mappings with much smaller end-to-end delay can be identified.

[36] *"Rigorous System Level Modeling and Analysis of Mixed HW/SW Systems".* Paraskevas Bourgos, Ananda Basu, Marius Bozga, Saddek Bensalem, Kai Huang and Joseph Sifakis. ACM/IEEE Ninth International Conference on Formal Methods and Models for Codesign (MEMOCODE'11). July 11th-13th, 2011. Cambridge, United Kingdom.

A grand challenge in complex embedded systems design is developing methods and tools for modeling and analyzing the behavior of an application software running on multicore or distributed platforms. We propose a rigorous method and a tool chain that allows to obtain a faithful model representing the behavior of a mixed hardware/software system from a model of its application software and a model of its underlying hardware architecture. The system model can be simulated and analyzed for validation of both functional and extra-functional properties. The tool chain uses DOL (Distributed Operation Layer, from ETH Zürich) as the frontend for specifying the application software and hardware architecture, and BIP (Behavior Interaction Priority, from VERIMAG) as the modeling and analysis framework. It is illustrated through the construction of system models of MJPEG

and MPEG2 decoder applications running on MPARM, a multicore architecture from Università di Bologna.

[37] *"Hierarchical Thermal Management Policy for High-Performance 3D Systems with Liquid Cooling"*. Francesco Zanini, Mohamed Sabry, David Atienza Alonso and Giovani De Micheli. IEEE Journal of Emerging and Selected Topics in Circuits and Systems (JETCAS), volume 1, issue 1, July 2011. DOI: 10.1109/JETCAS.2011.2158272. `http://infoscience.epfl.ch/record/165809`.

3-Dimensional integrated circuits and systems are expected to be present in electronic products in the short term. We consider specifically 3-D multi-processor systems-onchip (MPSoCs), realized by stacking silicon CMOS chips and interconnecting them by means of through-silicon vias (TSVs). Because of the high power density of devices and interconnect in the 3D stack, thermal issues pose critical challenges, such as hot-spot avoidance and thermal gradient reduction. Thermal management is achieved by a combination of active control of on-chip switching rates as well as active interlayer cooling with pressurized fluids. In this paper, we propose a novel online thermal management policy for high-performance 3D systems with liquid cooling. Our proposed controller uses a hierarchical approach with a global controller regulating the active cooling and local controllers (on each layer) performing dynamic voltage and frequency scaling (DVFS) and interacting with the global controller. Then, the online control is achieved by policies that are computed off-line by solving an optimization problem that considers the thermal profile of 3D-MPSoCs, its evolution over time and current time-varying workload requirements. The proposed hierarchical scheme is scalable to complex (and heterogeneous) 3D chip stacks. We perform experiments on a 3D-MPSoC case study with different interlayer cooling structures, using benchmarks ranging from web-accessing to playing multimedia. Results show significant advantages in terms of energy savings that reaches values up to 50% versus state-of-the-art thermal control techniques for liquid cooling, and thermal balance with differences of less than 10°C per layer.

[38] *"Attaining Single-Chip, High-Performance Computing Through 3D Systems with Active Cooling"*. Ayse K. Coskun, David Atienza, Mohamed Sabry and Jie Meng. IEEE Micro, vol. 16, num. 5, 2011, pages 25–35. `http://infoscience.epfl.ch/record/165525`.

Three-dimensional (3D) stacking is an attractive method for designing large manycore chips as it provides high transistor integration densities, improves manufacturing yield due to smaller chip area, reduces wirelength and capacitance, and enables heterogeneous integration of different technologies on the same chip. Stacking, however, significantly increases the thermal resistivity and the on-chip temperatures. In fact, temperature is among the major manufacturing challenges for 3D design. Active cooling, where the chip is cooled through the liquid flowing in built-in microchannels or through a cold plate, has emerged as a viable cooling alternative for high-performance 3D manycore systems. Liquid cooling is more efficient in removing the heat in comparison to conventional cooling methods with heat sinks and fans. Nevertheless, the dynamically changing nature of workloads running on manycore systems require runtime techniques to enable energy-efficient, reliable, and high-performance operation of liquid-cooled 3D systems. This article focuses on the benefits and the challenges of 3D design, and discusses novel techniques to integrate predictive cooling control with chip-level thermal management methods such as job scheduling and voltage frequency scaling. The key message is that 3D liquid-cooled systems with intelligent runtime management provide an energy-efficient solution to designing future single-chip high-performance manycore architectures.

[39] *"Real-Time Analysis of Servers for General Job Arrivals"*. Pratyush Kumar, Jian-Jia Chen, Lothar Thiele, Andreas Schranzhofer and Giorgio C. Buttazzo. Proceedings of the 17th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2011), August 28th-31st, 2011. Toyama, Japan.

Several servers have been proposed to schedule streams of aperiodic jobs in the presence of other periodic tasks. Standard schedulability analysis has been extended to consider such servers. However, not much attention has been laid on computing the worst-case delay suffered by a given stream of jobs when scheduled via a server. Such analysis is essential for using servers to schedule hard real-time tasks. We illustrate, with examples, that well established resource models, such as supply bound function and models from Real-Time Calculus, do not tightly characterize servers. In this work, we analyze the server algorithm of the Constant Bandwidth Server and compute a provably tight resource model of the server. The approach used enables us to differentiate between the soft and hard variants of the server. A similar approach can be used to characterize other servers; the final results for which are presented.

[40] *"System-Level Power and Timing Variability Characterization To Compute Thermal Guarantees"*. Pratyush Kumar and Lothar Thiele, CODES/ISSS 2011 at Embedded Systems Week, 9th-14th October, 2011. Taipei, Taiwan.

With ever-increasing power densities, temperature management using software and hardware techniques has become a necessity in the design of modern electronic systems. Such techniques have to be validated and optimized with respect to the thermal guarantee they provide, i.e., a safe upper-bound on the peak temperature of the system under all operating conditions. The computation of such a guarantee depends on the power and timing characteristics of the system. In this paper, we present formalisms to capture such characteristics at the system-level and provide an analytical technique to compute a provably safe upper-bound on the peak temperature. The proposed characterization and analysis is general in that it considers an impulse-response-based thermal model, task-dependent power consumption, tasks with dynamic arrival patterns and variable resource demand, and a scheduling policy expressed as a hierarchical composition of several commonly used policies.

[41] *"Compiling Applications for P2012 with the BIP Tool Chain"*. Ananda Basu, Marius Bozga, Sadddek Bensalem, Jean-Pierre Krimm, Julien Mottin, Christian Fabre, and François Pacull. 1st Platform 2012 Developper Conference, STMicroelectronics & CEA LETI, Minatec, Grenoble, France. December 2011.

This presentation describes the results of applying the rigorous system design ow based on the BIP framework on two embedded target applications. The first application is used within the SMECY project and is an image processing application for patterns and form recognition in aerospace and defence industries. The second application is also an image processing application dedicated to produce actual images from camera raw sensor output, and is used within PRO3D. We will detail step-by-step the engineering activities by which an embedded data-flow application written in plain C can be parallelized, then ported to and compiled by the BIP tool chain and finally be run on P2012 manycore simulators. For each application, we first provide a description in DOL. The corresponding system model in BIP is generated from the above description along with the description of the architecture and a mapping. Finally, we generate the deployable code from the system model which is executed on the P2012 simulators.

[42] *"PRO3D, an Introduction"*. Christian Fabre. October 2011, Software Technologies Concertation on Formal Methods for Components and Objects. October 3rd-5th, 2011, Torino, Italy. `http://fmco.liacs.nl/fmco11.html`.

This presentation provides an overview of PRO3D that tackles two 3D technologies and their consequences on stacked architectures and rigorous software development: through silicon vias (TSV) and liquid cooling. 3D memory hierarchies and the thermal impact of software on the 3D stack are explored. PRO3D experiments are mainly targeted on P2012, an industrial embedded manycore platform devlopped by STMicroelectronics.

[43] *"Low Cost Dynamic Voltage and Frequency Policy based upon Robust Control Techniques"*. Edith Beigné Sylvain Durand, Suzanne Lesecq and Christian Fabre. October 2011, Software Technologies Concertation on Formal Methods for Components and Objects. October 3rd-5th, 2011, Torino, Italy. `http://fmco.liacs.nl/fmco11.html`.
Today mobile computing platforms need ever-increasing computational performances while their energy consumption is drastically limited by battery lifespan. An optimal operating point is obtained thanks to a compromise between performance and power consumption. For distributed architectures (e.g. MultiProcessor System On Chip), the supply voltage and the operating frequency of each processing element are usually tuned dynamically to reach efficient performance/power consumption trade-offs. These control strategies (also called Dynamic Voltage and Frequency Scaling (DVFS)) can be implemented as closed-loop systems using control theory. This leads to the "sense and react" paradigm where information on the system state (especially actual temperature, occurrence of timing faults, supply voltage changes must first be gathered, possibly using data fusion techniques. Then from this information, and with robust and/or adaptive control approaches, the power consumption of the system can be finely controlled.

[44] *"Rigorous Component-Based System Design using the BIP Framework"*. Saddek Bensalem, Ananda Basu, Marius Bozga, Paraskevas Bourgos, and Joseph Sifakis. Invited presentation at the 5th Annual Layered Assurance Workshop, Orlando, Florida. December 2011.
.

[45] *"From BIP System Model to Platform P2012: A Code Generation Flow"*. Paraskevas Bourgos, Ananda Basu, Marius Bozga, and Saddek Bensalem. 1st Platform P2012 Developper Conference, Grenoble, France. December 2011.
We present an infrastructure for generating code for the P2012 platform from BIP system models. These models describe the behavior of mixed hardware/software systems. They can be simulated and formally verified using the BIP toolset. System models are obtained in a compositional and incremental manner, by source-to-source transformations, from descriptions of the application software, the hardware architecture, and the mapping. The above descriptions are given using the DOL framework. The generated code is targeted for the NPL runtime implemented for the P2012 platform, available within the 2011.1 SDK. This runtime provides API for thread management, memory allocation, communication and synchronization.

[46] *"Rigorous System Design: the BIP Approach"*. Ananda Basu, Saddek Bensalem, , Marius Bozga, Paraskevas Bourgos, and Joseph Sifakis. In *Proceedings of the Annual Doctoral Workshop on Mathematical and Engineering Methods in Computer Science, MEMICS 2011, Lednice, Czech Republic*, number 7119 in Lecture Notes in Computer Science. Springer, October 2011. Invited paper,
Rigorous system design requires the use of a single powerful component framework allowing the representation of the designed system at dierent levels of detail, from application software to its implementation. This is essential for ensuring the overall coherency and correctness. The paper introduces a rigorous design flow based on the BIP (Behavior, Interaction, Priority) component framework. This design flow relies on several, tool-supported, source-to-source transformations allowing to progressively and correctly transform high level application software towards efficient implementations for specific platforms.

[47] *"Automated Distributed Implementation of Component-based Models with Priorities"*. Borzoo Bonakdarpour, Marius Bozga, and Jean Quilbeuf. In *Proceedings of EMSOFT 2011, Taipei, Taiwan*, pages 59–68. ACM, October 2011.

In this paper, we introduce a novel model-based approach for constructing correct distributed implementation of component-based models constrained by priorities. We argue that model-based methods are especially of interest in the context of distributed embedded system due to their inherent complexity. Our three-phase method's input is a model specified in terms of a set of behavioural components that interact through a set of high-level synchronization primitives (e.g., rendezvous and broadcasts) and priority rules for scheduling purposes. Our technique, first, transforms the input model into a model that has no priorities. Then, it transforms the deprioritized model into another model that resolves distributed conflicts by incorporating a solution to the committee coordination problem. Finally, it generates distributed code using asynchronous point-to-point send/receive primitives. All transformations preserve the properties of their input model by ensuring observational equivalence. The transformations are implemented and our experiments validate their effectiveness.

[48] *"Algorithms for Synthesizing Priorities in Component-Based Systems"*. Chih-Hong Cheng, Saddek Bensalem, Yu-Fang Chen, Rongjie Yan, Barbara Jobstmann, Harald Ruess, Christian Buckl, and Alois Knoll. In *Proceedings of Automated Technology for Verification and Analysis, 9th International Symposium, ATVA 2011, Taipei, Taiwan*, volume 6996 of *Lecture Notes in Computer Science*, pages 150–167. Springer, October 2011.
We present algorithms to synthesize component-based systems that are safe and deadlock-free using priorities, which define stateless-precedence between enabled actions. Our core method combines the concept of fault- localization (using safety-game) and fault-repair (using SAT for conflict resolution). For complex systems, we propose three complementary methods as preprocessing steps for priority synthesis, namely (a) data abstraction to reduce component complexities, (b) alphabet abstraction and -deadlock to ignore components, and (c) automated assumption learning for compositional priority synthesis.

[49] *"Component Assemblies in the context of Manycore"*. Ananda Basu, Saddek Bensalem, Marius Bozga, Paraskevas Burgos, and Joseph Sifakis. Presentation at Software Technologies Concertation on Formal Methods for Components and Objects, FMCO 2011, Torino, Italy. October 2011.
.

[50] *"D-Finder 2: Towards Efficient Correctness of Incremental Design"*. Saddek Bensalem, Andreas Griesmayer, Axel Legay, Thanh-Hung Nguyen, Joseph Sifakis, and Rongjie Yan. In *Proceedings of NASA Formal Methods - Third International Symposium, NFM 2011, Pasadena, CA, USA*, volume 6617 of *Lecture Notes in Computer Science*, pages 453–458. Springer, April 2011.
D-Finder 2 is a new tool for deadlock detection in concurrent systems based on effective invariant computation to approximate the effects of interactions among modules. It is part of the BIP framework, which provides various tools centered on a component-based language for incremental design. The presented tool shares its theoretical roots with a previous implementation, but was completely rewritten to take advantage of a new version of BIP and various new results on the theory of invariant computation. The improvements are demonstrated by comparison with previous work and reports on new results on a practical case study.

[51] *"Model Construction and Priority Synthesis for Simple Interaction Systems"*. Chih-Hong Cheng, Saddek Bensalem, Barbara Jobstmann, Rongjie Yan, Alois Knoll, and Harald Ruess. In *Proceedngs of NASA Formal Methods - Third International Symposium, NFM 2011, Pasadena, CA, USA*, volume 6617 of *Lecture Notes in Computer Science*, pages 466–471. Springer, April 2011.
VissBIP is a software tool for visualizing and automatically orchestrating component-based systems consisting of a set of components and their possible interactions. The graphical interface of VissBIP

allows the user to interactively construct BIP models [3], from which executable code (C/C++) is generated. The main contribution of VissBIP is an analysis and synthesis engine for orchestrating components. Given a set of BIP components together with their possible interactions and a safety property, the VissBIP synthesis engine restricts the set of possible interactions in order to rule out unsafe states. The synthesis engine of VissBIP is based on automata-based (game-theoretic) notions. It checks if the system satisfies a given safety property. If the check fails, the tool automatically generates additional constraints on the interactions that ensure the desired property. The generated constraints define priorities between interactions and are therefore well-suited for conflict resolution between components.

[52] *"Design Space Exploration for 3D-Stacked DRAMs"*. C. Weis, N. Wehn, L. Igor, and L. Benini. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2011*, pages 1 –6, March 2011.
3D integration based on TSV (through silicon via) technology enables stacking of multiple memory layers and has the advantage of higher bandwidth at lower energy consumption for the memory interface. As in mobile applications energy efficiency is key, 3D integration is especially here a strategic technology. In this paper we focus on the design space exploration of 3D-stacked DRAMs with respect to performance, energy and area efficiency for densities from 256Mbit to 4Gbit per 3D-DRAM channel. We investigate four different technology nodes from 75nm down to 45nm and show the optimal design point for the currently most common commodity DRAM density of 1Gbit. Multiple channels can be combined for main memory sizes of up to 32GB. We present a functional SystemC model for the 3D-stacked DRAM which is coupled with a SDR/DDR 3D-DRAM channel controller. Parameters for this model were derived from detailed circuit level simulations. The exploration demonstrates that an optimized 1Gbit 3D-DRAM stack is 15× more energy efficient compared to a commodity Low-Power DDR SDRAM part without IO drivers and pads. To the best of our knowledge this is the first design space exploration for 3D-stacked DRAM considering different technologies and real world physical commodity DRAM data.

[53] *"Energy-Efficient Multi-Objective Thermal Control for Liquid-Cooled 3D Stacked Architectures"*. Mohamed Mostafa Sabri Aly, Ayse Kivilcim Coskun, David Atienza Alonso, Tajana Simunic Rosing, and Thomas Brunschwiler. *IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems*, 30(12):1883–1896, 2011.
3D stacked systems reduce communication delay in multiprocessor system-on-chips (MPSoCs) and enable heterogeneous integration of cores, memories, sensors, and RF devices. However, vertical integration of layers exacerbates temperatureinduced problems such as reliability degradation. Liquid cooling is a highly efficient solution to overcome the accelerated thermal problems in 3D architectures; however, it brings new challenges in modeling and run-time management for such 3D MPSoCs with multi-tier liquid cooling. This paper proposes a novel design-time/run-time thermal management strategy. The design-time phase involves a rigorous thermal impact analysis of various thermal control variables. We then utilize this analysis to design a run-time fuzzy controller for improving energy efficiency in 3D MPSoCs through liquid cooling management and dynamic voltage and frequency scaling (DVFS). The fuzzy controller adjusts the liquid flow rate dynamically to match the cooling demand of the chip for preventing over-cooling and for maintaining a stable thermal profile. The DVFS decisions increase chip-level energy savings and help balance the temperature across the system. Our controller is used in conjunction with temperatureaware load balancing and dynamic power management strategies. Experimental results on 2- and 4-tier 3D MPSoCs show that our strategy prevents the system from exceeding the given threshold temperature. At the same time, we reduce cooling energy by up to 63% and system-level energy by up to 21% in comparison to statically setting a flow rate setting to handle worst-case temperatures.

[54] *"Attaining Single-Chip, High-Performance Computing Through 3D Systems with Active Cooling"*. Ayse Kivilcim Coskun, David Atienza Alonso, Mohamed Mostafa Sabry Aly, and Jie Meng. *IEEE Micro*, 31(4):63–73, 2011.
Three-dimensional (3D) stacking is an attractive method for designing large manycore chips as it provides high transistor integration densities, improves manufacturing yield due to smaller chip area, reduces wirelength and capacitance, and enables heterogeneous integration of different technologies on the same chip. Stacking, however, significantly increases the thermal resistivity and the on-chip temperatures. In fact, temperature is among the major manufacturing challenges for 3D design. Active cooling, where the chip is cooled through the liquid flowing in built-in microchannels or through a cold plate, has emerged as a viable cooling alternative for high-performance 3D manycore systems. Liquid cooling is more efficient in removing the heat in comparison to conventional cooling methods with heat sinks and fans. Nevertheless, the dynamically changing nature of workloads running on manycore systems require runtime techniques to enable energy-efficient, reliable, and high-performance operation of liquid-cooled 3D systems. This article focuses on the benefits and the challenges of 3D design, and discusses novel techniques to integrate predictive cooling control with chip-level thermal management methods such as job scheduling and voltage frequency scaling. The key message is that 3D liquid-cooled systems with intelligent runtime management provide an energy-efficient solution to designing future single-chip high-performance manycore architectures.

[55] *"System-Level Thermal-Aware Design of 3D Multiprocessors with Inter-Tier Liquid Cooling"*. Arvind Sridhar, Mohamed M. Sabry, and David Atienza Alonso. In *Proceedings of Therminic 2011*, pages 1–9, Angleterre, 2011. EDA.
Rising chip temperatures and aggravated thermal reliability issues have characterized the emergence of 3D multiprocessor system-on-chips (3D-MPSoCs), necessitating the development of advanced cooling technologies. Microchannel based inter-tier liquid cooling of ICs has been envisaged as the most promising solution to this problem. A system-level thermal-aware design of electronic systems becomes imperative with the advent of these new cooling technologies, in order to preserve the reliable functioning of these ICs and effective management of the rising energy budgets of high-performance computing systems. This paper reviews the recent advances in the area of systemlevel thermal modeling and management techniques for 3D multiprocessors with advanced liquid cooling. These concepts are combined to present a vision of a green data center of the future which reduces the CO2 emissions by reusing the heat it generates.

[56] *"3D Thermal-Aware Floorplanner for Many-Core Single-Chip Systems"*. David Cuesta, Jose L. Risco, José Luis Ayala Rodrigo, and David Atienza Alonso. In *Proceedings of the 12th IEEE Latin-American Test Workshop (LATW11)*, volume 1, pages 1–6, New York, 2011. IEEE Press.
Heat removal and power density distribution delivery have become two major reliability concerns in 3D stacked technology. In this paper, we propose a thermal-driven 3D floorplanner. Our contributions include: (1) a novel multi-objective formulation to consider the thermal and performance constraints in the optimization approach; (2) an efficient Mixed Integer Linear Programming (MILP) representation of the floorplanning model; and (3) a smooth integration of the MILP model with an accurate thermal modelling of the architecture. The experimental work is conducted for two realistic many-core single-chip architectures: an homogeneous system resembling Intel's SCC, and an improved heterogeneous setup. The results show promising improvements of the mean, peak temperature and the thermal gradient, with a reduced overhead in the wire length of the system.

[57] *"Towards Thermally-Aware Design of 3D MPSoCs with Inter-Tier Cooling"*. Mohamed Sabry, Arvind Sridhar, David Atienza Alonso, Yuksel Temiz, Yusuf Leblebici, Sylwia Szczukiewicz, Navid Borhani, John Richard Thome, Thomas Brunschwiler, and Bruno

Michel. In *Proceedings of Design, Automation and Test in Europe (DATE '11)*, pages 1466–1471. ACM and IEEE Press, 2011.

New tendencies envisage 3D Multi-Processor System-On-Chip (MPSoC) design as a promising solution to keep increasing the performance of the next-generation highperformance computing (HPC) systems. However, as the power density of HPC systems increases with the arrival of 3D MPSoCs, supplying electrical power to the computing equipment and constantly removing the generated heat is rapidly becoming the dominant cost in any HPC facility. Thus, both power and thermal/cooling implications play a major role in the design of new HPC systems, given the energy constraints in our society. Therefore, EPFL, IBM and ETHZ have been working within the CMOSAIC Nano-Tera.ch program project in the last three years on the development of a holistic thermally-aware design. This paper presents the exploration in CMOSAIC of novel cooling technologies, as well as suitable thermal modeling and system-level design methods, which are all necessary to develop 3D MPSoCs with inter-tier liquid cooling systems. As a result, we develop energy-efficient run-time thermal control strategies to achieve energy-efficient cooling mechanisms to compress almost 1 Tera nano sized functional units into one cubic centimeter with a 10 to 100 fold higher connectivity than otherwise possible. The proposed thermally-aware design paradigm includes exploring the synergies of hardware-, software- and mechanical-based thermal control techniques as a fundamental step to design 3D MPSoCs for HPC systems. More precisely, we target the use of inter-tier coolants ranging from liquid water and twophase refrigerants to novel engineered environmentally friendly nano-fluids, as well as using specifically designed micro-channel arrangements, in combination with the use of dynamic thermal management at system-level to tune the flow rate of the coolant in each micro-channel to achieve thermally-balanced 3D-ICs. Our management strategy prevents the system from surpassing the given threshold temperature while achieving up to 67% reduction in cooling energy and up to 30% reduction in system-level energy in comparison to setting the flow rate at the maximum value to handle the worst-case temperature.

[58] *"Thermal Analysis and Active Cooling Management for 3D MPSoCs"*. Mohamed Sabry, David Atienza Alonso, and Ayse Kivilcim Coskun. In *Proceedings of IEEE International Symposium on Circuits and Systems (ISCAS'11)*, volume 1, pages 2237–2240, New York, 2011. IEEE Press.

3D stacked architectures reduce communication delay in multiprocessor system-on-chips (MPSoCs) and allowing more functionality per unit area. However, vertical integration of layers exacerbates the reliability and thermal problems, and cooling is a limiting factor in multi-tier systems. Liquid cooling is a highly efficient solution to overcome the accelerated thermal problems in 3D architectures. However, liquid cooling brings new challenges in modeling and run-time management. This paper proposes a design-time/run-time thermal management policy for 3D MPSoCs with inter-tier liquid cooling. First, we perform a design-time analysis to estimate the thermal impact of liquid cooling and dynamic voltage frequency scaling (DVFS) on 3D MPSoCs. Based on this analysis, we define a set of management rules for run-time thermal management. We utilize these rules to control and adjust the liquid flow rate in order to match the cooling demand for preventing energy wastage of overcooling, while maintaining a stable thermal profile in the 3D MPSoCs. Experimental results on multi-tier 3D MPSoCs show that proposed design-time/run-time management policy prevents the system to exceed the given threshold temperature while reducing cooling energy by 50% on average and system-level energy by 18% on average in comparison to using a static worstcase flow rate setting.

[59] *"Thermal-Aware System-Level Design of 2D/3D MPSoC Architectures"*. David Atienza Alonso. Keynote at the 23rd Euromicro Conference on Real-Time Systems (ECRTS 2011). July 2011, Porto, Portugal. `http://www.cister.isep.ipp.pt/ecrts11/ECRTS11-keynote-Atienza.pdf`.

Multi-Processor Systems-on-Chips (MPSoCs) are penetrating the consumer electronics market as powerful solutions to the growing demand for scalable and high-performance systems, at limited design complexity and power dissipation. Nevertheless, MPSoCs are prone to alarming temperature variations on the die, which seriously decrease their expected reliability and lifetime. Furthermore, technical advances in manufacturing technologies are fueling the trend towards high performance 3D MPSoC designs. However, 3D stacking creates higher power and heat density, leading to further degrading reliability and performance if thermal management is not handled properly. Thus, it is critical to develop dedicated design methodologies that guarantee safe thermal behavior of forthcoming 2D and 3D MPSoCs at low energy and performance cost. In this talk I discuss the development of novel system-level design methodologies for 2D and 3D MPSoCs that seamlessly address thermal modeling, analysis and management. First, I will revise thermal modeling mechanisms for 2D MPSoCs based on simulation and emulation frameworks. Second, I will introduce reactive and proactive run-time thermal management methods which prevent hot spots and large thermal gradients in 2D MPSoCs while incurring negligible performance degradation. Finally, I will show how new thermal modeling and active management methods, including liquid cooling, can be modeled and included in this novel design flows for next-generation 3D MPSoC architectures.

[60] *"System-Level Thermal-Aware Design of 3D Multi-Processors with Inter-Tier Liquid Cooling"*. David Atienza Alonso. Keynote at the 17th International Workshop on Thermal investigations of ICs and Systems (THERMINIC 2011). September 2011, Paris, France. `http://cmp.imag.fr/conferences/therminic/therminic2011/Invitedspeakers.php`.
Continuous advances in manufacturing technologies are enabling the development of more powerful and compact 3D multi-processor ICs for high-performance computing (HPC). However, 3D stacking originates higher power and heat densities, leading to degraded performance in HPC if cooling is not handled properly at all levels of abstraction. In this talk, I present a novel thermal-aware design paradigm for 3D-ICs and overall "Green Datacenter Design", developed at EPFL in cooperation with IBM and ETHZ, which includes thermal modeling as fundamental step to design 3D multi-processor ICs with inter-tier liquid cooling microchannels, in combination with the use of dynamic thermal management at system-level to tune the flow rate of the coolant in each tier in order to achieve thermally-balanced 3D-ICs and complete datacenters servers.

[61] *"Supporting OpenMP on a Multi-Cluster Embedded MPSoC"*. Andrea Marongiu, Paolo Burgio, and Luca Benini. *Microprocessors and Microsystems*, 35(8):668–682, 2011. Design and Verification of Complex Digital Systems. `http://www.sciencedirect.com/science/article/pii/S0141933111001001`,
The ever-increasing complexity of MPSoCs is putting the production of software on the critical path in embedded system development. Several programming models and tools have been proposed in the recent past that aim to facilitate application development for embedded MPSoCs. OpenMP is a mature and easy-to-use standard for shared memory programming, which has recently been successfully adopted in embedded MPSoC programming as well. To achieve performance, however, it is necessary that the implementation of OpenMP constructs efficiently exploits the many peculiarities of MPSoC hardware, and that custom features are provided to the programmer to control it. In this paper we consider a representative template of a modern multi-cluster embedded MPSoC and present an extensive evaluation of the cost associated with supporting OpenMP on such a machine, investigating several implementation variants that are aware of the memory hierarchy and of the heterogeneous interconnection.

[62] *"MPOpt-Cell: a high-performance data-flow programming environment for the CELL BE processor"*. Alessio Franceschelli, Paolo Burgio, Giuseppe Tagliavini, Andrea Marongiu,

Martino Ruggiero, Michele Lombardi, Alessio Bonfietti, Michela Milano, and Luca Benini. In *Proceedings of the 8th ACM International Conference on Computing Frontiers*, CF '11, pages 11:1–11:2, New York, NY, USA, May 2011. ACM.
We present MPOpt-Cell, an architecture-aware framework for high-productivity development and efficient execution of stream applications on the CELL BE Processor. It enables developers to quickly build Synchronous Data Flow (SDF) applications using a simple and intuitive programming inter- face based on a set of compiler directives that capture the key abstractions of SDF. The compiler backend and system runtime efficiently manage hardware resources.

[63] *"A distributed and topology-agnostic approach for on-line NoC testing"*. M.R. Kakoee, V. Bertacco, and L. Benini. In *Networks on Chip (NoCS), 2011 Fifth IEEE/ACM International Symposium on*, pages 113 –120, May 2011.
A new distributed on-line test mechanism for NoCs is proposed which scales to large-scale networks with general topologies and routing algorithms. Each router and its links are tested using neighbors in different phases. Only the router under test is in test mode and all other parts of the NoC are in functional mode. Experimental results show that our on-line test approach can detect stuck-at and short-wire faults in the routers and links. Our approach achieves 100 % fault coverage for the data-path and 85 % for the control paths including routing logic, FIFO's control path and the arbiter of a 5x5 router. Synthesis results show that the hardware overhead of our test components with TMR (Triple Module Redundancy) support is 20 % for covering both stuck-at and short-wire faults and 7 % for covering only stuck-at faults in the 5x5 router. Simulation results show that our on-line testing approach has an average latency overhead of 20 % and 3 % in synthetic traffic and PARSEC traffic benchmarks on an 8x8 NoC, respectively.

[64] *"Static Thermal Model Learning for High-Performance Multicore Servers"*. F. Beneventi, A. Bartolini, and L. Benini. In *Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN)*, pages 1 –6, July 2011.
Aggressive thermal management is a critical feature for high-end computing platforms, as worst-case thermal budgeting is becoming unaffordable. Reactive thermal management, which sets temperature thresholds to trigger thermal capping actions, is too "near-sighted", and it may lead to severe performance degradation and thermal overshoots. More aggressive proactive thermal management minimizes performance penalty with smooth optimal control, but it requires the knowledge of the system thermal models to be precise. Unfortunately, in practice these models are not provided by equipment manufacturers, and they strongly depend on the deployment environment. Hence, we need to develop procedures to derive thermal models automatically in the field. In this paper, we focus on static thermal model learning. We tackle the problem in a real-life context: we developed a complete infrastructure for model-building and thermal data collection in the Linux environment, and we tested it on an Intel Nehalem-based server CPU. Model building is based on a least-square procedure which extracts the model linking power dissipation with temperature in steady-state conditions. Our results show high accuracy and robustness even in presence of a complex thermal environment and limited-precision power and temperature measurements typical of today's commercial servers.

[65] *"A Distributed and Self-Calibrating Model-Predictive Controller for Energy and Thermal Management of High-Performance Multicores"*. A. Bartolini, M. Cacciari, A. Tilli, and L. Benini. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2011*, pages 1 –6, March 2011.
High-end multicore processors are characterized by high power density with significant spatial and temporal variability. This leads to power and temperature hot-spots, which may cause non-uniform ageing and accelerated chip failure. These critical issues can be tackled on-line by closed-loop ther-

mal and reliability management policies. Model predictive controllers (MPC) outperform classic feedback controllers since they are capable of minimizing a cost function while enforcing safe working temperature. Unfortunately basic MPC controllers rely on a-priori knowledge of multicore thermal model and their complexity exponentially grows with the number of controlled cores. In this paper we present a scalable, fully-distributed, energy-aware thermal management solution. The model-predictive controller complexity is drastically reduced by splitting it in a set of simpler interacting controllers, each allocated to a core in the system. Locally, each node selects the optimal frequency to meet temperature constraints while minimizing the performance penalty and system energy. Global optimality is achieved by letting controllers exchange a limited amount of information at run-time on a neighbourhood basis. We address model uncertainty by supporting learning of the thermal model with a novel distributed self-calibration approach that matches well the controller architecture.

[66] *"Variability-Aware Task Allocation for Energy-Efficient Quality of Service Provisioning in Embedded Streaming Multimedia Applications"*. F. Paterna, A. Acquaviva, A. Caprara, F. Papariello, G. Desoli, and L. Benini. *Computers, IEEE Transactions on*, PP(99):1, 2011.

Multimedia streaming applications running on next-generation parallel multiprocessor arrays in sub-45nm technology face new challenges related to device and process variability, leading to performance and power variations across the cores. In this context, Quality of Service (QoS), as well as energy efficiency, could be severely impacted by variability. In this work we propose a run-time variability-aware workload distribution technique for enhancing real-time predictability and energy efficiency based on an innovative Linear-Programming + Bin-Packing formulation which can be solved in linear time. We demonstrate our approach on the virtual prototype of a next-generation industrial multi-core platform running a multithread MPEG2 decoder. Experimental results confirm that our technique compensates variability, while improving energy-efficiency and minimizing deadline violations in presence of performance and power variations across the cores.

[67] *"A fully-synthesizable single-cycle interconnection network for Shared-L1 processor clusters"*. A. Rahimi, I. Loi, M.R. Kakoee, and L. Benini. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2011*, pages 1 –6, March 2011.

Shared L1 memory is an interesting architectural option for building tightly-coupled multi-core processor clusters. We designed a parametric, fully combinational Mesh-of-Trees (MoT) interconnection network to support high-performance, single-cycle communication between processors and memories in L1-coupled processor clusters. Our interconnect IP is described in synthesizable RTL and it is coupled with a design automation strategy mixing advanced synthesis and physical optimization to achieve optimal delay, power, area (DPA) under a wide range of design constraints. We explore DPA for a large set of network configurations in 65nm technology. Post placement&routing delay is 38FO4 for a configuration with 8 processors and 16 32-bit memories (8×16); when the number of both processors and memories is increased by a factor of 4, the delay increases almost logarithmically, to 84FO4, confirming scalability across a significant range of configurations. DPA tradeoff flexibility is also promising: in comparison to the maxperformance 16×32 configuration, there is potential to save power and area by 45 % and 12 % respectively, at the expense of 30 % performance degradation.

[68] *"A System Level Approach to Multi-core Thermal Sensors Calibration"*. Andrea Bartolini, MohammadSadegh Sadri, Francesco Beneventi, Matteo Cacciari, Andrea Tilli, and Luca Benini. In José Ayala, Braulio García-Cámara, Manuel Prieto, Martino Ruggiero, and Gilles Sicard, editors, *Integrated Circuit and System Design. Power and Timing Modeling, Optimization, and Simulation*, volume 6951 of *Lecture Notes in Computer Science*, pages

22–31. Springer Berlin / Heidelberg, 2011.
Many-cores systems on chip provide the highest performance scaling potential due to the massive parallelism, but they suffer from thermal issues due to their high power densities. Thermal sensors and feedback strategies are used to reduce these threats but sensor accuracy directly impact control performance. In this paper we propose a novel technique to calibrate thermal sensors. Our approach can be applied to general multi-core platforms since it combines stress patterns and least-square fitting to perform thermal sensor characterization directly on the target device. We experimentally validate our approach on the Single Chip Cloud (SCC) prototype by Intel.

[69] *"Single-Chip Cloud Computer thermal model"*. M. Sadri, A. Bartolini, and L. Benini. In *Thermal Investigations of ICs and Systems (THERMINIC), 2011 17th International Workshop on*, pages 1 –6, September 2011.
Spatial and temporal non-uniformities of workload and power consumption advanced Systems-on-Chip (SoC) platforms result in localized high power densities, which lead to temperature hot-spots, gradients and thermal cycles that may cause non-uniform ageing and accelerated chip failure. The Single-Chip Cloud Computer (SCC) is an experimental many-core processor created by Intel Labs and it integrates thermal sensors to track the chip thermal behavior. Unfortunately these sensors provide a limited introspection on the full-chip thermal map, as they monitor temperature at a coarse granularity. In this paper we build a fine-grained thermal and power model of SCC using a state-of-the-art thermal modeling tool (Hotspot). We calibrate the model with measured data from chip sensors. We assess the predictive power of the thermal model and its main sources of error.

[70] *"Exploring Instruction Caching Strategies for Tightly-Coupled Shared-Memory Clusters"*. D. Bortolotti, F. Paterna, C. Pinto, A. Marongiu, M. Ruggiero, and L. Benini. In *System on Chip (SoC), 2011 International Symposium on*, pages 34 –41, October 2011.
Several Chip-Multiprocessor designs today leverage tightly-coupled computing clusters as a building block. These clusters consist of a fairly large number N of simple cores, featuring fast communication through a shared multibanked L1 data memory and ≈ 1 Instruction-Per-Cycle (IPC) per core. Thus, aggregated I-fetch bandwidth approaches $f * N$, where f is the cluster clock frequency. An effective instruction cache architecture is key to support this I-fetch bandwidth. In this paper we compare two main architectures for instruction caching targeting tightly coupled CMP clusters: (i) private instruction caches per core and (ii) shared instruction cache per cluster. We developed a cycle-accurate model of the tightly coupled cluster with several configurable architectural parameters for exploration, plus a programming environment targeted at efficient data-parallel computing. We conduct an in-depth study of the two architectural templates based on the use of both synthetic microbenchmarks and real program workloads. Our results provide useful insights and guidelines for designers.

[71] *"Fast and Lightweight Support for Nested OpenMP Parallelism on a Multi-Cluster Platform"*. Andrea Marongiu, Paolo Burgio, and Luca Benini. Presented at the 1st Plateform 2012 Developper Conference, Grenoble, France. December 2011.
OpenMP is a well-known programming model for shared memory parallelism. It consists of a set of compiler directives, library routines and environment variables that provide a simple means to specify parallel execution within a sequential code. OpenMP was originally designed for Symmetric Multi-Processors (SMP), but recently many implementations for embedded MPSoCs have been proposed. In this paper we describe an implementation of the OpenMP runtime library for a multi-cluster MPSoC, modeled after the P2012 architectural template.

[72] *"Instruction Cache Architectures for Tightly-Coupled Shared-Memory Clusters"*. Daniele Bortolotti, Francesco Paterna, Christian Pinto, Andrea Marongiu, Martino Ruggiero, and Luca Benini. Presented at the 1st Plateform 2012 Developper Conference, Grenoble,

France. December 2011.
To keep the pace of Moore's law, several Chip-Multiprocessors (CMP) platforms are embracing the many-core paradigm, where a large number of simple cores are integrated onto the same die. Current examples of many-cores include GP-GPUs such as NVIDIA Fermi, the HyperCore Architecture Line (HAL) processors from Plurality, or STMicroelectronics' Platform 2012. All of the cited architectures share a few common traits: their fundamental computing tile is a tightly coupled multicore cluster with a shared multibanked L1 memory for fast data access and a fairly large number of simple cores. Key to providing instruction-fetch bandwidth for a cluster is an effective instruction cache architecture design. The main contribution of this work is the analysis and comparison of the two main architectures for instruction caching targeting tightly coupled CMP clusters: (1) private instruction caches per core and (2) shared instruction cache per cluster. We developed a cycle-accurate model of a P2012-like cluster with the two cache organizations, and with several configurable architectural parameters for exploration.

[73] *"OpenMP-based HW Acceleration for On-Chip Multi-Core Shared-Memory Clusters"*. Andrea Marongiu, Paolo Burgio, and Luca Benini. Presented at the 1st Plateform 2012 Developper Conference, Grenoble, France. December 2011.
Key to designing accelerator-based MPSoCs in a cost-effective manner is the availability of methodologies to quickly define and instantiate accelerators within a suitable architec- tural template, from both a hardware and a software per- spective. By clearly defining such templates, and providing streamlined communication and synchronization mechanisms between processors and accelerators, programming models can be enriched with abstract constructs to allow designers to focus on accelerator specification at a high level. We present a vertically integrated HW/SW architecture, with a programming model and runtime support to design tightly coupled clusters including one or more dedicated accelerators named HW Processing Units (HWPU). The proposed approach includes an extended OpenMP programming API and compiler that allows the designer to mix code parallelization and acceleration mechanisms while hiding implementation details. Specifically, we extend OpenMP with a key custom directive to outline code regions which are to be hardware-accelerated, rather than executed in software.

[74] *"Verification of the P2012 Manycore Architecture: Underlying Issues"*. Richard Hersemeule. October 2011, Software Technologies Concertation on Formal Methods for Components and Objects. October 3rd-5th, 2011, Torino, Italy. `http://fmco.liacs.nl/fmco11.html`.
This presentation provides an overview of hardware verification issues for the P2012 manycore platform.

[75] *"Mapping of Applications to MPSoCs"*. Peter Marwedel, Juergen Teich, Georgia Kouveli, Iuliana Bacivarov, Lothar Thiele, Soonhoi Ha, Chanhee Lee, Qiang Xu, and Lin Huang. CODES+ISSS 201111, Taipei, Taiwan, October 2011.
The advent of embedded many-core architectures results in the need to come up with techniques for mapping embedded applications onto such architectures. This paper presents a representative set of such techniques. The techniques focus on optimizing performance, temperature distribution, reliability and fault tolerance for various models.

Scientific Contributions Published During the 3rd Reporting Period (2012)

[76] *"PRO3D: Programming for Future 3D Manycore Architectures"*. Christian Fabre. January 2012, Invited talk at the 6th "Interconnection Network Architecture: On-Chip, Multi-Chip

(INA-OCMC)" held in conjunction with the 7th. HiPEAC Conference, January 25, 2012, Paris, France.
Presentation of [77].

[77] *"PRO3D: Programming for Future 3D Manycore Architectures"*. Christian Fabre, Iuliana Bacivarov, Lothar Thiele, Hoeseok Yang, Pratyush Kumar, Devesh Chokshi, Ahmed Jerraya, Julien Mottin, Jean-Pierre Krimm, Ananda Basu, Saddek Bensalem, Marius Bozga, Paraskevas Bourgos, Martino Ruggiero, Luca Benini, Andrea Marongiu, Eric Flamand, and Diego Melpignano. In *Proceeddings of the 6th Workshop on Interconnection Network Architecture: On-Chip, Multi-Chip (INA-OCMC 2012)*, Paris, France, January 2012. ACM Digital Library.
PRO3D tackles two 3D technologies and their consequences on stacked architectures and software stack: through silicon vias (TSV) and liquid cooling. 3D memory hierarchies and the thermal impact of software on the 3D stack are mainly explored. The PRO3D software development flow is based on a rigorous assembly of software components and monitors the thermal integrity of the 3D stack. PRO3D experiments are mainly targeted on P2012, an industrial embedded manycore platform devlopped by STMicroelectronics.

[78] *"Power Mode Selection in Embedded Systems with Performance Constraints"*. Y. Akgul, D. Puschini, S. Lesecq, I. Miro-Panades, P. Benoit, L. Torres, and E. Beigné. In *Faible Tension Faible Consommation (FTFC), 2012 IEEE*, pages 1 –4, June 2012.
Mobile computing platforms must provide ever increasing performances under stringent power consumption constraints. Dynamic Voltage and Frequency Scaling (DVFS) techniques allow to reduce the power consumption by providing just enough power to the chip in order to finish the task before its deadline. DVFS is usually achieved by setting the supply voltage and the clock frequency to predefined values (so-called "power modes") during given durations that depend on the task to be run and on its deadline. Here, the problem of power management is recast as a linear programming one and the computation of the duration spent in each one of the N power modes is obtained with a Simplex algorithm solution. Results for 3 power modes exemplify the proposed approach.

[79] *"VT-State Condition Monitoring in Integrated Circuits Using Fusion of Information From General Purpose Sensors"*. S. Lesecq, L. Vincent, E. Beigné, and Ph Maurines. In *12th International Forum on Embedded MPSoC and Multicore, July 9th-13rd, 2012, Québec, Canada*, July 2012. `http://www.mpsoc-forum.org/previous/2012/index.htm`,
Today mobile computing platforms need ever-increasing computational performances while their energy consumption is drastically limited by battery lifespan. An optimal operating point is obtained thanks to a compromise between performance and power consumption. For distributed architectures, the supply voltage and the operating frequency of each processing element can be tuned dynamically to reach efficient performance/power consumption trade-offs. As a consequence, the physical state (e.g. actual supply voltage and temperature) of the integrated circuit must be monitored to locally adapt the chip parameters. A new method has been developed to estimate the supply voltage and temperature of a local area in an integrated circuit. The raw measurements are acquired form standard ring oscillators buried in the chip and the sensor fusion technique makes use of statistical tests.

[80] *"Local (V,T) Estimation in Integrated Circuits Using a Set of Statistical Tests"*. S. Lesecq, L. Vincent, E. Beigné, and Ph Maurines. Keynote at the IEEE International Conference on Advanced Technologies for Communications (ATC/REV), Hanoi, Vietnam, October 10–12, 2012. October 2012.

Mobile platforms need ever-increasing computational performances under stringent energy consumption limitation mainly due to the battery lifespan. An optimal operating point is obtained thanks to a compromise between performance and power consumption. For distributed architectures (e.g. MultiProcessor System on Chip), the supply voltage and the operating frequency of each processing element are dynamically tuned to reach efficient performance/power consumption tradeoffs. To increase the performance of each "actuator", the physical state (e.g. the current supply voltage and temperature) of the integrated circuit must be monitored to locally adapt the control parameters. During this keynote, we will present a new estimation method based on statistical tests to estimate the supply voltage and the temperature of a local area in an integrated circuit. Standard ring oscillators buried in the chip provide the raw measurements that are fused to estimate the IC physical state.

[81] *"Advanced coupled voltage-frequency control for power efficient DVFS management"*. Carolina Albea-Sanchez, Diego Puschini, Suzanne Lesecq, and Yeter Akgul. In *38th Annual Conference of the IEEE Industrial Electronics Society*, page 99, Montréal, Canada, October 2012.
Power management is a hot-topic in complex System-on-Chip (SoC) designs. In the context of advanced technologies, Dynamic Voltage-Frequency Scaling (DVFS) techniques are widely proposed to improve e ciency. Nowadays, these mechanisms are composed of independent actuators controlling the applied voltage and clock frequency. A prede ned sequence has to be used to switch from one state to another in order to avoid timing faults but increasing the energy cost. The timing of the sequence depends on the dynamic response of actuators. In this work, an external controller is designed in order to couple both actuators to manage the voltage and frequency transient periods, increasing power e ciency The proposed controller has been implemented to couple a Vdd-hopping mechanism with a Frequency-Lock Loop circuit.

[82] *"How State Estimation in Integrated Circuits Based on Statistical Tests Can Be Used to Fine-Tune the Control of the Voltage and Frequency Actuators in the Power Management Framework"*. S. Lesecq, L. Vincent, E. Beigné, and Ph Maurines. Keynote at VARI 2012, the 3rd European Workshop on CMOS Variability, Nice, France. June 2012.
Today mobile computing platforms need ever-increasing computational performances while their energy consumption is drastically limited by battery lifespan. An optimal operating point is obtained thanks to a compromise between performance and power consumption. For distributed architectures (e.g. MultiProcessor System on Chip), the supply voltage and the operating frequency of each processing element are dynamically tuned to reach efficient performance/power consumption tradeoffs. To increase the performance of each "actuator", the physical state (e.g. its current supply voltage and temperature) of the integrated circuit must be monitored to locally adapt the control parameters. During this keynote, we will present a new estimation method based on statistical tests to estimate the supply voltage and the temperature of a local area in an integrated circuit. The raw measurements are acquired form standard ring oscillators buried in the chip and they are fused to estimate the IC physical state. Then we will show how this information might be used to fine tune the control of the closed-loop actuators in order to ensure for these actuators the appropriate functioning, whatever the physical state in a given range.

[83] *"Knowledge-Based Distributed Conflict Resolution for Multiparty Interactions and Priorities"*. Saddek Bensalem, Marius Bozga, Jean Quilbeuf, and Joseph Sifakis. In Holger Giese and Grigore Rosu, editors, *FMOODS/FORTE - Formal Techniques for Distributed Systems - Joint 14th IFIP WG 6.1 International Conference, FMOODS 2012 and 32nd IFIP WG 6.1 International Conference, FORTE 2012, Stockholm, Sweden, June 13-16, 2012. Proceedings*, volume 7273 of *Lecture Notes in Computer Science*, pages 118–134.

Springer, June 2012.
Distributed decentralized implementation of systems of communicating processes raises non-trivial problems. Correct execution of multiparty interactions, subject to priority rules, requires sophisticated mechanisms for runtime conflict detection and resolution. We propose a method for detection of false conflicts which combines partial observation of the system's state and apriori knowledge extracted from invariants. We propose heuristics for determining optimal sets of observations leading to implementations with particular guarantees. We provide preliminary experimental results on an implementation of the method in the BIP framework – See poster on Figure 3 page 45.

[84] *"Rigorous Component-Based System Design"*. Saddek Bensalem, Ananda Basu, Marius Bozga, Paraskevas Bourgos, and Joseph Sifakis. In Francisco Duran, editor, *9th International Workshop on Rewriting Logic and its Applications WRLA 2012, Tallinn, Estonia March 24-25, 2012. Pre proceedings*, pages 1–6. Institute of Cybernetics at Tallinn University of Technology, March 2012.
Rigorous system design requires the use of a single powerful component framework allowing the representation of the designed system at different levels of detail, from application software to its implementa- tion. This is essential for ensuring the overall coherency and correctness. The paper introduces a rigorous design flow based on the BIP (Behavior, Interaction, Priority) component framework [1]. This design flow relies on several, tool-supported, source-to-source transformations allowing to progressively and correctly transform high level application software to- wards efficient implementations for specific platforms.

[85] *"A framework for automated distributed implementation of component-based models"*. Borzoo Bonakdarpour, Marius Bozga, Mohamad Jaber, Jean Quilbeuf, and Joseph Sifakis. *Distributed Computing*, 25(5):383–409, October 2012.
Although distributed systems are widely used nowadays, their implementation and deployment are still time-consuming, error-prone, and hardly predictable tasks. In this paper, we propose a method for producing automatically efficient and correct-by-construction distributed implementations from a model of the application software in Behavior, Interaction, Priority (BIP). BIP is a well-founded component-based framework encompassing high-level multi-party interactions for synchronizing components (e.g., rendezvous and broadcast) and dynamic priorities for scheduling between interactions. Our method transforms an arbitrary BIP model into a Send/Receive BIP model that is directly implementable on distributed execution platforms. The transformation consists in (1) breaking the atomicity of actions in components by replacing synchronous multiparty interactions with asynchronous Send/Receive interactions; (2) inserting distributed controllers that coordinate the execution of interactions according to a user-defined partition of interactions, and (3) adding a distributed algorithm for handling conflicts between controllers. The obtained Send/Receive BIP model is proven observationally equivalent to its corresponding initial model. Hence, all functional properties of the initial BIP model are preserved by construction in the implementation. Moreover, the obtained Send/Receive BIP model can be used to automatically derive distributed executable code. The proposed method is fully implemented. Currently, it is possible to generate C++ implementations for (1) TCP sockets for conventional distributed communication, (2) MPI for multiprocessor platforms, and (3) POSIX threads for deployment on multi-core platforms. We present four case studies and report experimental results for different design choices including partition of interactions and choice of algorithm for distributed conflict resolution.

[86] *"Model-based implementation of distributed systems with priorities"*. Borzoo Bonakdarpour, Marius Bozga, and Jean Quilbeuf. *Design Automation for Embedded Systems*, pages 1–26, July 2012.
Model-based application development aims at increasing the application's integrity by using models

employed in clearly defined transformation steps leading to correct-by-construction artifacts. In this paper, we introduce a novel model-based approach for constructing correct distributed implementation of component-based models constrained by priorities. We argue that model-based methods are especially of interest in the context of distributed embedded systems due to their inherent complexity (e.g., caused by non-deterministic nature of distributed systems). Our method is designed based on three phases of transformation. The input is a model specified in terms of a set of behavioral components that interact through a set of high-level synchronization primitives (e.g., rendezvous and broadcasts) and priority rules for scheduling purposes. The first phase transforms the input model into a model that has no priorities. Then, the second phase transforms the deprioritized model into another model that resolves distributed conflicts by incorporating a solution to the committee coordination problem. Finally, the third phase generates distributed code using asynchronous point-to-point message passing primitives (e.g., TCP sockets). All transformations preserve the properties of their input model by ensuring observational equivalence. All the transformations are implemented and our experiments validate their effectiveness.

[87] *"Statistical Model Checking QoS Properties of Systems with SBIP"*. Saddek Bensalem, Marius Bozga, Benoît Delahaye, Cyrille Jégourel, Axel Legay, and Ayoub Nouri. In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods, Verification and Validation. Technologies for Mastering Change - 5th International Symposium, ISoLA 2012, Proceedings, Part I*, volume 7609 of *Lecture Notes in Computer Science*, pages 327–341. Springer, 2012.
BIP is a component-based framework supporting rigorous design of embedded systems. This paper presents SBIP, an extension of BIP that relies on a new stochastic semantics that enables verication of large-size systems by using Statistical Model Checking. The approach is illustrated on several industrial case studies.

[88] *"Optimized Distributed Implementation of Multiparty Interactions with Observation"*. Saddek Bensalem, Marius Bozga, Jean Quibeuf, and Joseph Sifakis. In *2nd International Workshop on Programming based on Actors, Agents, and Decentralized Control Workshop held at ACM SPLASH 2012, Pre-Proceedings*, pages 90–99, October 2012.
Using high level coordination primitives allows enhanced expressiveness of component-based frameworks to cope with the inherent complexity of present-day systems designs. Nonetheless, their distributed implementation raises multiple issues, regarding both the correctness and the runtime performance of the final implementation. We propose a novel approach for distributed implementation of multiparty interactions subject to scheduling constraints expressed by priorities.We rely on new composition operators and semantics that combine multiparty interactions with observation. We show that this model provides a natural encoding for priorities and moreover, can be used as an intermediate step towards provably correct and optimized distributed implementations.

[89] *"Knowledge Based Transactional Behavior"*. Saddek Bensalem, Marius Bozga, Jean Quilbeuf, and Doron Peled. In *Eight Haifa Verification Conference, HVC2012*, November 2012. Post proceedings to appear as LNCS volume,
Component-based systems (including distributed programs and multiagent systems) involve a lot of coordination. This coordination is done in the background, and is transparent to the operation of the system. The reason for this overhead is the interplay between concurrency and nondeterministic choice: processes alternate between progressing independently and coordinating with other processes, where coordination can involve multiple choices of the participating components. This kind of interactions appeared as early as some of the main communication-based programming languages, where overhead effort often causes a restriction on the possible coordination. With the goal of enhancing the efficiency of coordination for component-based systems, we propose here a

method for coordination-based on the precalculation of the knowledge of processes and coordination agents. This knowledge can be used to lift part of the communication or synchronization that appears in the background of the execution to support the interaction. Our knowledge-based method is orthogonal to the actual algorithms or primitives that are used to guarantee the synchronization: it only removes messages conveying information that knowledge can infer.

[90] *"Neural Network-Based Thermal Simulation of Integrated Circuits on GPUs"*. Arvind Sridhar, Alessandro Vincenzi, Martino Ruggiero, and David Atienza Alonso. *IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems*, 31(1):23–36, 2012.
With the rising challenges in heat removal in integrated circuits (ICs), the development of thermal-aware computing architectures and run-time management systems have become indispensable to the continuation of IC design scaling. These thermal-aware design technologies of the future strongly depend on the availability of efficient and accurate means for thermal modeling and analysis. These thermal models must have not only the sufficient accuracy to capture the complex mechanisms that regulate thermal diffusion in ICs, but also a level of abstraction that allows for their fast execution for design space exploration. In this paper, we propose an innovative thermal modeling approach for full-chips that can handle the scalability problem of transient heat flow simulation in large 2D/3D multi-processor ICs. This is achieved by parallelizing the computation-intensive task of transient temperature tracking using neural networks and exploiting the computational power of massively parallel graphics processing units (GPUs). Our results show up to 35x run-time speed-up compared to state-of-the-art IC thermal simulation tools while keeping the error lower than $1^{o}C$. Speed-ups scale with the size of the 3D multi-processor ICs and our proposed method serves as a valuable design space exploration tool.

[91] *"Thermal Balancing of Liquid-Cooled 3D-MPSoCs Using Channel Modulation"*. Mohamed M. Sabry, Arvind Sridhar, and David Atienza Alonso. In *EDAA 2012*, 2012.
While possessing the potential to replace conventional air-cooled heat sinks, inter-tier microchannel liquid cooling of 3D ICs also creates the problem of increased thermal gradients from the fluid inlet to outlet ports [1, 2]. These cooling-induced thermal gradients can be high enough to create undesirable stress in the ICs, undermining the structural reliability and lifetimes. In this paper, we present a novel design-time solution for the thermal gradient problem in liquid-cooled 3D Multi-Processor System-on-Chip (MPSoC) architectures. The proposed method is based on channel width modulation and provides the designers with an additional dimension in the design-space exploration. We formulate the channel width modulation as an optimal control design problem to minimize the temperature gradients in the 3D IC while meeting the design constraints. The proposed thermal balancing technique uses an analytical model for forced convective heat transfer in microchannels, and has been applied to a two tier 3D-MPSoC. The results show that the proposed approach can reduce thermal gradients by up to 31% when applied to realistic 3D-MPSoC architectures, while maintaining pressure drops in the microchannels well below their safe limits of operation.

[92] *"GreenCool: An Energy-Efficient Liquid Cooling Design Technique for 3D MPSoCs Via Channel Width Modulation"*. Mohamed Mostafa Sabry Aly, Arvind Sridhar, Jie Meng, Ayse Kivilcim Coskun, and David Atienza Alonso. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 31(12):1–14, 2012.
Liquid cooling using interlayer microchannels has appeared as a viable and scalable packaging technology for 3D multiprocessor system-on-chips (MPSoCs). Microchannelbased liquid cooling, however, can substantially increase the onchip thermal gradients, which are undesirable for reliability, performance, and cooling efficiency. In this work we present GreenCool, an optimal design methodology for liquid-cooled 3D MPSoCs. GreenCool simultaneously minimizes the cooling energy for a given system while maintaining thermal gradients and peak temperatures under safe

limits. This is accomplished by tuning the heat transfer characteristics of the microchannels using channel width modulation. Channel width modulation is compatible with the current process technologies and incurs minimal additional fabrication costs. Through an extensive set of experiments, we show that channel width modulation is capable of complementing and enhancing the benefits of temperature-aware floorplanning. We also experiment with a 16-core 3D system with stacked DRAM, for which GreenCool improves energy efficiency by up to 53% compared to cooling optimization without channel modulation.

[93] *"An energy efficient DRAM subsystem for 3D integrated SoCs"*. C. Weis, I. Loi, L. Benini, and N. Wehn. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2012*, pages 1138 –1141, march 2012.
Energy efficiency is the key driver for the design optimization of System-on-Chips for mobile terminals (smartphones and tablets). 3D integration of heterogeneous dies based on TSV (through silicon via) technology enables stacking of multiple memory or logic layers and has the advantage of higher bandwidth at lower energy consumption for the memory interface. In this work we propose a highly energy efficient DRAM subsystem for next-generation 3D integrated SoCs, which will consist of a SDR/DDR 3D-DRAM controller and an attached 3D-DRAM cube with a fine-grained access and a very flexible (WIDE-IO) interface. We implemented a synthesizable model of the SDR/DDR 3D-DRAM channel controller and a functional model of the 3D-stacked DRAM which embeds an accurate power estimation engine. We investigated different DRAM families (WIDE IO DDR/SDR, LPDDR and LPDDR2) and densities that range from 256Mb to 4Gb per channel. The implementation results of the proposed 3D-DRAM subsystem show that energy optimized accesses to the 3D-DRAM enable an overall average of 37% power savings as compared to standard accesses. To the best of our knowledge this is the first design of a 3D-DRAM channel controller and 3D-DRAM model featuring co-optimization of memory and controller architecture.

[94] *"Quantifying the Impact of Frequency Scaling on the Energy Efficiency of the Single-Chip Cloud Computer"*. A. Bartolini, M. Sadri, J. Furst, A.K. Coskun, and L. Benini. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2012*, pages 181 –186, march 2012.
Dynamic frequency and voltage scaling (DVFS) techniques have been widely used for meeting energy constraints. Single-chip many-core systems bring new challenges owing to the large number of operating points and the shift to message passing interface (MPI) from shared memory communication. DVFS, however, has been mostly studied on single-chip systems with one or few cores, without considering the impact of the communication among cores. This paper evaluates the impact of frequency scaling on the performance and power of many-core systems with MPI. We conduct experiments on the Single-Chip Cloud Computer (SCC), an experimental many-core processor developed by Intel. The paper first introduces the run-time monitoring infrastructure and the application suite we have designed for an in-depth evaluation of the SCC. We provide an extensive analysis quantifying the effects of frequency perturbations on performance and energy efficiency. Experimental results show that run-time communication patterns lead to significant differences in power/performance tradeoffs in many-core systems with MPI.

[95] *"A Distributed Interleaving Scheme for Efficient Access to WideIO DRAM Memory"*. Ciprian Seiculescu, Luca Benini, and Giovanni De Micheli. In *Proceedings of the Eighth IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis*, CODES+ISSS '12, pages 103–112, New York, NY, USA, October 2012. ACM.
Achieving the main memory (DRAM) required bandwidth at acceptable power levels for current and future applications is a major challenge for System-on-Chip designers for mobile platforms. Three dimensional (3D) integration and 3D stacked DRAM memories promise to provide a significant

boost in bandwidth at low power levels by exploiting multiple channels and wide data interfaces. In this paper, we address the problem of efficiently exploiting the multiple channels provided by standard (JEDEC's WIDE-IO) 3D-stacked memories, to extract maximal effective bandwidth and minimize latency for main memory access. We propose a new distributed interleaved access method that leverages the on-chip interconnect to simplify the design and implementation of the DRAM controller, without impacting performance compared to traditional centralized implementations. We perform experiments on realistic workload for a mobile communication and multimedia platform and show that our proposed distributed interleaving memory access method improves the overall throughput while minimally impacting the performance of latency sensitive communication flows.

[96] *"Full system simulation of many-core heterogeneous SoCs using GPU and QEMU semi-hosting"*. Shivani Raghav, Andrea Marongiu, Christian Pinto, David Atienza, Martino Ruggiero, and Luca Benini. In *Proceeding of the 5th Annual Workshop on General Purpose Processing with Graphics Processing Units GPGPU-5*, pages 101–109. ACM, May 2012.

Modern system-on-chips are evolving towards complex and heterogeneous platforms with general purpose processors coupled with massively parallel manycore accelerator fabrics (e.g. embedded GPUs). Platform developers are looking for efficient full-system simulators capable of simulating complex applications, middleware and operating systems on these heterogeneous targets. Unfortunately current virtual platforms are not able to tackle the complexity and heterogeneity of state-of-the-art SoCs. Software emulators, such as the open-source QEMU project, cope quite well in terms of simulation speed and functional accuracy with homogeneous coarse-grained multi-cores. The main contribution of this paper is the introduction of a novel virtual prototyping technique which exploits the heterogeneous accelerators available in commodity PCs to tackle the heterogeneity challenge in full-SoC system simulation. In a nutshell, our approach makes it possible to partition simulation between the host CPU and GPU. More specifically, QEMU runs on the host CPU and the simulation of manycore accelerators is offloaded, through semi-hosting, to the host GPU. Our experimental results confirm the flexibility and efficiency of our enhanced QEMU environment.

[97] *"An OpenMP Compiler for Efficient Use of Distributed Scratchpad Memory in MPSoCs"*. A. Marongiu and L. Benini. *Computers, IEEE Transactions on*, 61(2):222 –236, February 2012.

Most of today's state-of-the-art processors for mobile and embedded systems feature on-chip scratchpad memories. To efficiently exploit the advantages of low-latency high-bandwidth memory modules in the hierarchy, there is the need for programming models and/or language features that expose such architectural details. On the other hand, effectively exploiting the limited on-chip memory space requires the programmer to devise an efficient partitioning and distributed placement of shared data at the application level. In this paper, we propose a programming framework that combines the ease of use of OpenMP with simple, yet powerful, language extensions to trigger array data partitioning. Our compiler exploits profiled information on array access count to automatically generate data allocation schemes optimized for locality of references.

[98] *"SIMinG-1k: A Thousand-Core Simulator Running on General-Purpose Graphical Processing Units"*. Shivani Raghav, Andrea Marongiu, Christian Pinto, Martino Ruggiero, David Atienza, and Luca Benini. *Concurrency and Computation: Practice and Experience*, pages n/a–n/a, October 2012. `http://dx.doi.org/10.1002/cpe.2940`,

This paper introduces SIMinG-1k – a manycore simulator infrastructure. SIMinG-1k is a graphics processing unit accelerated, parallel simulator for design-space exploration of large-scale manycore systems. It features an optimal trade-off between modeling accuracy and simulation speed. Its main objectives are high performance, flexibility, and ability to simulate thousands of cores. SIMinG-1k

can model different architectures (currently, we support ARM and Intel x86) using two-step approach where architecture specific front end is decoupled from a fast and parallel manycore virtual machine running on graphical processing unit platform. We evaluate the simulator for target architecture with up to 4096 cores. Our results demonstrate very high scalability and almost linear speedup with simulation of increasing number of cores.

[99] *"P2012: Building an ecosystem for a scalable, modular and high-efficiency embedded computing accelerator"*. Luca Benini, Eric Flamand, Didier Fuin, and Diego Melpignano. In Rosenstiel and Thiele [101], pages 983–987.
P2012 is an area- and power-efficient many-core computing fabric based on multiple globally asynchronous, locally synchronous (GALS) clusters supporting aggressive fine-grained power, reliability and variability management. Clusters feature up to 16 processors and one control processor with independent instruction streams sharing a multi-banked L1 data memory, a multi-channel DMA engine, and specialized hardware for synchronization and scheduling. P2012 achieves extreme area and energy efficiency by supporting domain-specific acceleration at the processor and cluster level through the addition of dedicated HW IPs. P2012 can run standard OpenCL and OpenMP parallel codes well as proprietary Native Programming Model (NPM) SW components that provide the highest level of control on application-to-resource mapping. In Q3 2011 the P2012 SW Development Kit (SDK) has been made available to a community of R&D users; it includes full OpenCL and NPM development environments. The first P2012 SoC prototype in 28 *nm CMOS will sample in Q4 2012, featuring four clusters and delivering 80GOPS (with single precision floating point support) in* 15.2 mm^2 *with* 2 *W power consumption.*

[100] *"Platform 2012, a many-core computing accelerator for embedded SoCs: performance evaluation of visual analytics applications"*. Diego Melpignano, Luca Benini, Eric Flamand, Bruno Jego, Thierry Lepley, Germain Haugou, Fabien Clermidy, and Denis Dutoit. In Groeneveld et al. [102], pages 1137–1142.
P2012 is an area- and power-efficient many-core computing accelerator based on multiple globally asynchronous, locally synchronous processor clusters. Each cluster features up to 16 processors with independent instruction streams sharing a multi-banked one-cycle access L1 data memory, a multi-channel DMA engine and specialized hardware for synchronization and aggressive power management. P2012 is 3D stacking ready and can be customized to achieve extreme area and energy efficiency by adding domain-specific HW IPs to the cluster. The first P2012 SoC prototype in 28 *nm CMOS will sample in Q3, featuring four 16-processor clusters, a 1MB L2 memory and delivering 80GOPS (with 32 bit single precision floating point support) in* 18 mm^2 *with* 2 *W power consumption (worst-case). P2012 can run standard OpenCL™ and proprietary Native Programming Model SW components to achieve the highest level of control on application-to-resource mapping. A dedicated version of the OpenCV vision library is provided in the P2012 SW Development Kit to enable visual analytics acceleration. This paper will discuss preliminary performance measurements of common feature extraction and tracking algorithms, parallelized on P2012, versus sequential execution on ARM CPUs.*

[101] *2012 Design, Automation & Test in Europe Conference & Exhibition, DATE 2012, Dresden, Germany, March 12-16, 2012*. Wolfgang Rosenstiel and Lothar Thiele, editors. IEEE, March 2012.
(Includes PRO3D publications.).

[102] *The 49th Annual Design Automation Conference 2012, DAC '12, San Francisco, CA, USA, June 3-7, 2012*. Patrick Groeneveld, Donatella Sciuto, and Soha Hassoun, editors. ACM, June 2012.
(Includes PRO3D publications.).

[103] *"Platform 2012 Many-Core Programmable Accelerator: Status and Perspectives"*. Eric Flamand and Diego Melpignano. March 2012, DATE 2012, Dresden, Germany, March 12-16, 2012.
Invited Speaker during the session "Many-Core Architectures and Compilers" – Organizer: Maurizio Palesi.

[104] *"Timing Analysis on a Processor with Temperature-Controlled Speed Scaling"*. Pratyush Kumar and Lothar Thiele. In *In Proc. of the 18th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2012*, Beijing, China, April 2012.
Several recent works consider the problem of temperature-constrained scheduling of jobs. In such attempts, speed of the processor and the execution of jobs is software controlled such that temperature and performance constraints are met. An alternative approach is to use measurements from temperature sensors to actuate the speed of the processor as a feedback control loop. Though such a solution explicitly and independently meets the thermal constraints, the analysis of the real-time properties of tasks served by such a processor is not straightforward. In this paper, we study this problem for a variable stream of jobs characterized by an input arrival rate. We show that an intuitive notion of monotonicity extends to such a processor. Using this property, we present an analytical technique to determine the worst-case delay suffered by jobs. The presented technique efficiently and tightly determines the delay as a function of the initial temperature. The simplicity of this analysis motivates further analysis and mainstream use of such systems.

[105] *"Worst-Case Temperature Guarantees for Real-Time Applications on Multi-Core Systems"*. Lars Schor, Iuliana Bacivarov, Hoeseok Yang, and Lothar Thiele. In *In Proc. of the 18th IEEE Real-Time and Embedded Technology and Applications Symposium, RTAS 2012*, Beijing, China, April 2012. IEEE Computer.
Due to increased on-chip power density, multi-core systems face various thermal issues. In particular, exceeding a certain threshold temperature can reduce the system's performance and reliability. Therefore, when designing a real-time application with non-deterministic workload, the designer has to be aware of the maximum possible temperature of the system. This paper proposes an analytic method to calculate an upper bound on the worst-case peak temperature of a realtime system with multiple cores generated under all possible scenarios of task executions. In order to handle a broad range of uncertainties, task arrivals are modeled as periodic event streams with jitter and delay. Finally, the proposed method is applied to a multi-core ARM platform and our results are validated in various case studies.

[106] *"Fast Worst-Case Peak Temperature Evaluation for Real-Time Applications on Multi-Core Systems"*. Lars Schor, Iuliana Bacivarov, Hoeseok Yang, and Lothar Thiele. In *Proc. IEEE Latin American Test Workshop (LATW)*, Quito, Ecuador, April 2012. IEEE.
The reliability of multi-core systems is nowadays threatened by high chip temperatures leading to long-term reliability concerns and short-term functional errors. In real-time systems, high chip temperatures are even adherent to potential deadline violations. Therefore, correct functionality can only be guaranteed if the worst-case peak temperature is incorporated in real-time analysis. However, calculating the peak temperature of hundreds of design alternatives during design space exploration is time-consuming. In this paper, we address this challenge and present a fast analytic method to calculate a non-trivial upper bound on the maximum temperature of a multi-core realtime system with non-deterministic workload. The considered thermal model is able to address various thermal effects like heat exchange between neighboring cores and temperature-dependent leakage power. Finally, the proposed method is applied to a multicore ARM platform to validate its efficiency and accuracy.

[107] *"A Hybrid Approach to Cyber-Physical Systems Verification"*. Pratyush Kumar, Dip Goswami, Dipanjan Chakraborty, Kai Lampka, Anuradha Annaswamy, and Lothar Thiele. In *In Proceedings of the 49th Design Automation Conference, DAC 2012*, San Fransisco, USA, June 2012. ACM.

We propose a performance verification technique for cyber-physical systems that consist of multiple control loops implemented on a distributed architecture. The architectures we consider are fairly genric and arise in domains such as automotive and industrial automation; there are multiple processor or electronic control unites (ECUs) communicating over buses like FlexRay and CAN. Current practice involves analyzing the architecture to estimate worst-case end-to-end message delays and using these delays to design the control applications. This involves a significant amount of pessimism since the worst-case delays often occur very rarely. We show how to combine functional analysis techniques with model checking in order to derive a delay-frequency interface that quantifies the interleavings between messages with worst-case delays and those with smaller delays. In other words, we bound the frequency with which control messages might suffer the worst-case delay. We show that such a delay-frequency interface enables us to verify much tighter control performance properties compares to the what would be possibly only with worst-case bounds.

[108] *"Quantifying the Effect of Rare Timing Events with Settling-Time and Overshoot"*. Pratyush Kumar and Lothar Thiele. In *In Proceedings of the 33rd Real-Time Systems Symposium, RTSS 2012*, San Juan, Puerto Rico, December 2012. IEEE.

For hard real-time systems, worst-case timing models are employed to validate whether timeliness properties, such as meeting deadlines, are always satisfied. We argue that such a deadline-interface should be generalised in view of two separate motivations: (a) applications can tolerate bounded non-satisfaction of timeliness properties due to inherent ro- bustness or relaxed quality requirements, and (b) worst-case timing models do not expose the occurrence of certain rare yet predictable events. As a more expressive interface, we propose the Rare-Event with Settling-Time (REST) model wherein, during rare events nominal timing models can be violated up to a known bound. Such a violation may lead to non-satisfaction of the timeliness properties up to a certain bound. We characterise this bound in terms of (a) the longest interval when the deadlines are not met, which we call the settling-time, and (b) the maximum number of jobs that can miss deadlines during the settling-time called the overshoot. We propose two models of rare events, characterised on an interval domain. For a single stream of jobs, we provide methods to tightly compute the settling-time and overshoot. For multiple streams of jobs on a single processor, we show that amongst schedulers agnostic to the occurrence of the rare event, the EDF scheduler optimally minimises the settling-time. In contrast, RM is not optimal within the class of fixed priority schedulers.

Scientific Contributions Published After the Project (2013)

[109] *"PRO3D, Programming for Future 3D Manycore Architectures: Project Interim Status"*. Christian Fabre, Iuliana Bacivarov, Ananda Basu, Martino Ruggiero, David Atienza, Eric Flamand, Jean-Pierre Krimm, Julien Mottin, Lars Schor, Pratyush Kumar, Hoeseok Yang, Devesh B. Chokshi, Lothar Thiele, Saddek Bensalem, Marius Bozga, Luca Benini, Mohamed M. Sabry, Yusuf Leblebici, Giovanni De Micheli, and Diego Melpignano. In Beckert et al. [111]. To appear,

PRO3D tackles two important 3D technologies, that are Through Silicon Via (TSV) and liquid cooling, and investigates their consequences on stacked architectures and entire software development. In particular, memory hierarchies are being revisited and the thermal impact of software on the 3D stack is explored. As a key result, a software design flow based on the rigorous assembly of

software components and monitoring of the thermal integrity of the 3D stack has been developed. After 30 months of research, PRO3D proposes a complete tool-chain for 3D manycore, that integrates state-of-the-art tools ranging from system-level formal specification and 3D exploration, to actual programming and runtime control on the 3D system. Current efforts are directed towards extensive experiments on an industrial embedded manycore platform.

[110] *"Low-Cost Dynamic Voltage and Frequency Management based upon Robust Control Techniques under Thermal Constraints"*. Sylvain Durand, Suzanne Lesecq, Edith Beigné, Christian Fabre, Lionel Vincent, and Diego Puschini. In Beckert et al. [111], pages 334–353. To appear,

Mobile computing platforms need ever increasing performance, which implies an increase in the clock frequency applied to the processing elements (PE). As a consequence, the distribution of a single global clock over the whole circuit is tremendously difficult. Globally Asynchronous Locally Synchronous (GALS) designs alleviate the problem of clock distribution by having multiple clocks, each one being distributed on a small area of the chip. Energy consumption is the main limiting factor for mobile platforms as they are powered by batteries. Dynamic Voltage and Frequency Scaling (DVFS) in each Voltage and Frequency Island (VFI) has proven to be highly effective to reduce the power consumption of the chip while meeting the performance requirements. Environmental parameters (i.e. temperature and supply voltage) changes also strongly affect the chip performance and its power consumption. Some sensors can be buried in order to estimate via data fusion techniques the supply voltage and the temperature variations. For instance the knowledge of the gap between the temperature and its maximum value can be used to adapt the power management technique. The present paper deals with the design of a voltage and frequency management approach (DVFS) that explicitly takes into account the thermal constraints of the platform.

[111] *Formal Methods for Components and Objects, 10th International Symposium FMCO 2011. State-of-the-Art Survey*. Bernhard Beckert, Ferruccio Damiani, Frank de Boer, and Marcello Bonsangue, editors. volume 7542 of *LNCS*. Springer, 2013.
(Includes PRO3D articles.).

[112] *"Component Assemblies in the Context of Many-Core"*. Ananda Basu, Saddek Bensalem, Marius Bozga, Paraskevas Bourgos, Mayur Maheshwari, and Jospeh Sifakis. In Beckert et al. [111]. To appear,

We present a component-based software design flow for building parallel applications running on top of manycore platforms. The flow is based on the BIP - Behaviour, Interaction, Priority - component framework and its associated toolbox. It provides full support for modeling of application software, validation of its functional correctness, modeling and performance analysis on system-level models, code generation and deployment on target manycore platforms. The paper details some of the steps of the design flow. The design flow is illustrated through the modeling and deployment of two applications, the Cholesky factorization and the MJPEG decoding on MPARM, an ARM-based manycore platform. We emphasize the merits of the design flow, notably fast performance analysis as well as code generation and efficient deployment on manycore platforms.

[113] *"Online Thermal Control Methods for Multi-Processor Systems"*. Francesco Zanini, David Atienza Alonso, Colin Jones, Luca Benini, and Giovanni De Micheli. *ACM Transactions on Design Automation of Electronic Systems*, 18(1):1–24, 2013.

With technological advances, the number of cores integrated on a chip is increasing. This, in turn is leading to thermal constraints and thermal design challenges. Temperature gradients and hot-spots not only affect the performance of the system, but also lead to unreliable circuit operation and affect the life-time of the chip. Meeting temperature constraints and reducing hot-spots are critical for achieving reliable and efficient operation of complex multi-core systems. In this article

we analyze the use of four of the most promising families of online control techniques for thermal management of multi-processors system-on-chip (MPSoC). In particular, in our exploration we aim at achieving an online smooth thermal control action that minimizes the performance loss as well as the computational and hardware overhead of embedding a thermal management system inside the MPSoC. The definition of the optimization problem to tackle in this work considers the thermal profile of the system, its evolution over time and current time-varying workload requirements. Thus, this problem is formulated as a finite-horizon optimal control problem and we analyze the control features of different on-line thermal control approaches. In addition, we implemented the policies on an MPSoC hardware simulation platform and performed experiments on a cycle-accurate model of the 8-core Niagara multi-core architecture using benchmarks ranging from web-accessing to playing multimedia. Results show different trade-offs among the analyzed techniques regarding the thermal profile, the frequency setting, the power consumption and the implementation complexity.

[114] *"A Satisfiability Approach to Speed Assignment for Distributed Real-Time Systems"*. Pratyush Kumar, Devesh Chokshi, and Lothar Thiele. In *In Proceedings of the 2013 Design, Automation & Test in Europe Conference & Exhibition, DATE 2013*, Grenoble, France, March 2013. IEEE.
We study the problem of assigning speeds to resources serving distributed applications with delay, buffer and energy constraints. We argue that the considered problem does not have any straightforward solution due to the intricately related constraints. As a solution, we propose using Real-Time Calculus (RTC) to analyse the constraints and a SATisfiability solver to efficiently explore the design space. To this end, we develop an SMT solver by using the OpenSMT framework and the Modular Performance Analysis (MPA) toolbox. Two key enablers for this implementation are the analysis of incomplete models and generation of conflict clauses in RTC. The results on problem instances with very large decision spaces indicate that the proposed SMT solver performs very well in practice.

[115] *"Predictability for Timing and Temperature in Multiprocessor System-on-Chip Platforms"*. Lothar Thiele, Lars Schor, Iuliana Bacivarov, and Hoeseok Yang. *ACM Transactions in Embedded Computing Systems (TECS)*, March 2013.
High computational performance in multiprocessor system-on-chips (MPSoCs) is constrained by the ever-increasing power densities in integrated circuits, so that nowadays MPSoCs face various thermal issues. For instance, high chip temperatures may lead to long-term reliability concerns and short-term functional errors. Therefore, the new challenge in designing embedded real-time MPSoCs is to guarantee the final performance and correct function of the system, considering both functional and non-functional properties. One way to achieve this is by ruling out mapping alternatives that do not fulfill requirements on performance or peak temperature already in early design stages. In this article, we propose a thermal-aware optimization framework for mapping real-time applications onto MPSoC platforms. The performance and temperature of mapping candidates are evaluated by formal temporal and thermal analysis models. To this end, analysis models are automatically generated during design space exploration, based on the same specifications as used for software synthesis. The analysis models are automatically calibrated with performance data reflecting the execution of the system on the target platform. The data is automatically obtained prior to design space exploration based on a set of benchmark mappings. Case studies show that the performance and temperature requirements are often conflicting goals and optimizing them together leads to major benefits in terms of a guaranteed and predictable high performance.

[116] *"Thermal-Aware Task Assignment for Real-Time Applications on Multi-Core Systems"*. Lars Schor, Hoeseok Yang, Iuliana Bacivarov, and Lothar Thiele. In Beckert et al. [111], pages 294–313.

The reduced feature size of electronic systems and the demand for high performance lead to increased power densities and high chip temperatures, which in turn reduce the system reliability. Thermal-aware task allocation and scheduling algorithms are promising approaches to reduce the peak temperature of multi-core systems with real-time constraints. However, as long as the worst-case chip temperature is not incorporated into system analysis, no guarantees on the performance can be given. This paper explores thermal-aware task assignment strategies for real-time applications with non-deterministic workload that are running on a multi-core system. In particular, tasks are assigned to the multi-core system so that the worst-case chip temperature is minimized and all real-time deadlines are met. Each core has its own clock domain and the static assigned frequency corresponds to the minimum operation frequency such that no real-time deadline is missed. Finally, we show that the proposed temperature minimization problem can efficiently be solved by metaheuristics.

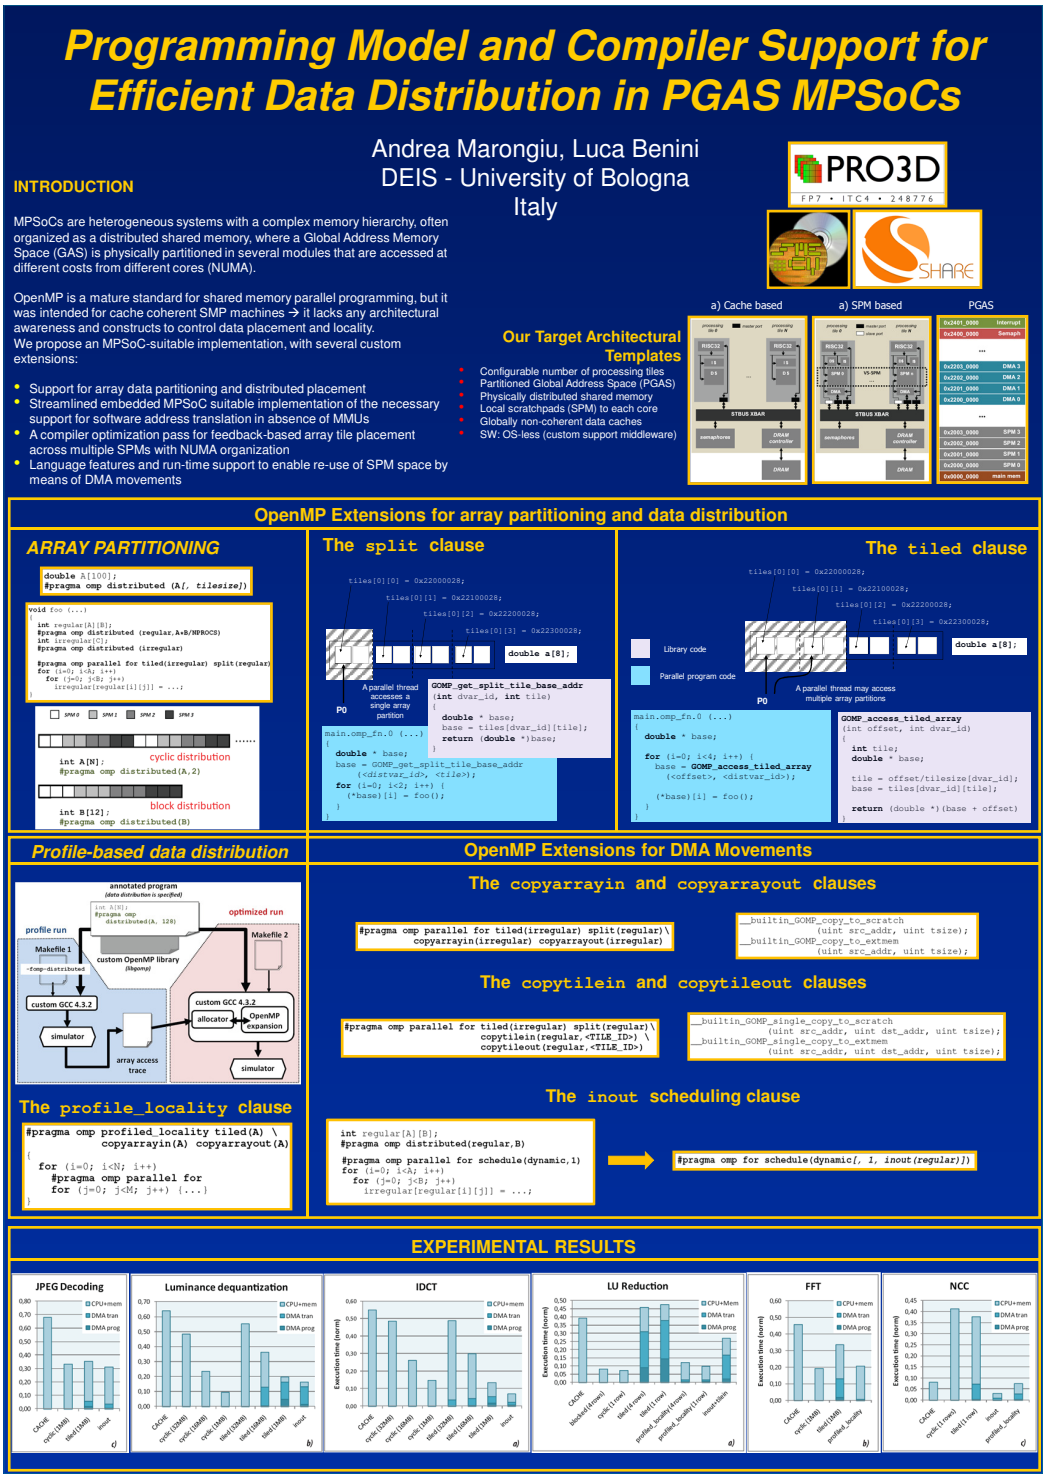

Figure 2: Poster Presented at Intel ERIC – Braunschweig, Germany – Sep. 19th-23rd, 2010

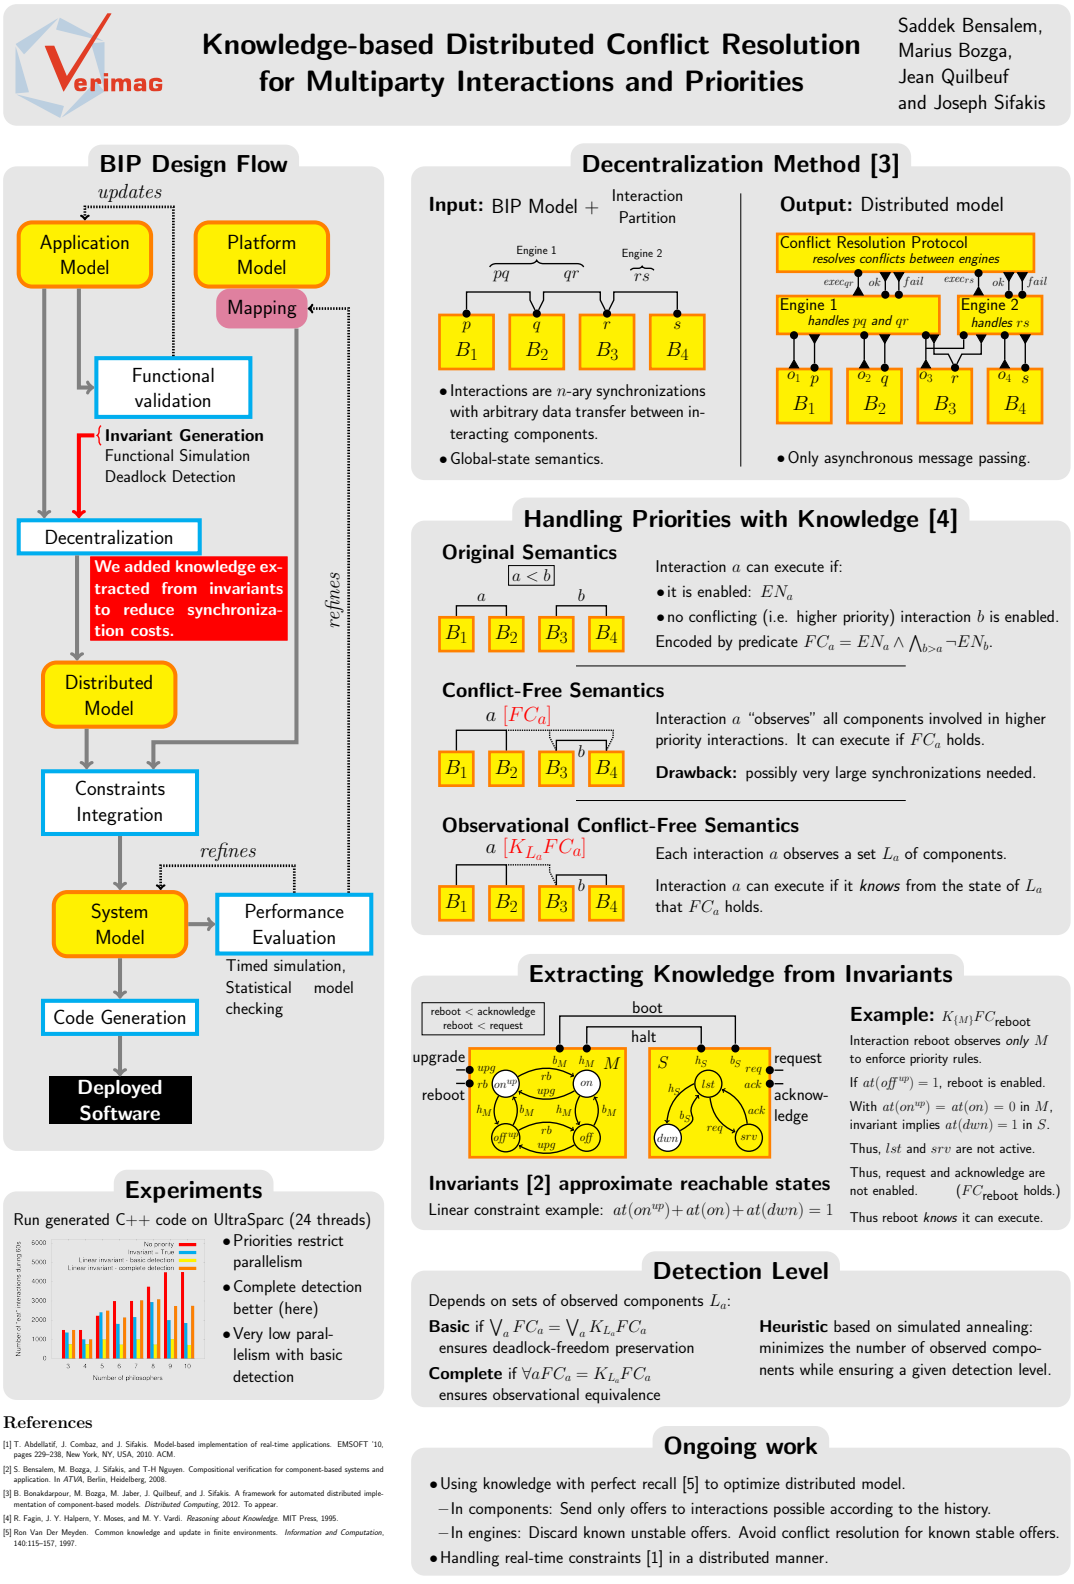

Figure 3: Poster Presented at FORTE 2012 – Stockholm, Sweden – June 13rd-16th, 2012