IST-004527 ARTIST2 NoE   Year 2
Cluster: Execution Platforms, Compilers & TA   D16-EP-Y2
Activity: Resource Aware Design (JPRA-NoE Integration)

ARTIST 2

Network of Excellence

IST-004527 ARTIST2:
Embedded Systems Design

Activity Progress Report for Year 2

JPRA-NoE Integration

Resource aware design

Clusters:

Execution Platforms
Compilers and Timing Analysis

Activity Leaders:

Prof. Dr. Peter Marwedel, University of Dortmund
http://ls12-www.cs.uni-dortmund.de/~marwedel

Prof. Luca Benini, University of Bologna
http://www-micrel.deis.unibo.it/~benini/

Policy Objective (abstract)

Provide, through the integration of research activities of many participants, a viable path for resource-aware software and hardware development. The final objective is to achieve integration of research activities in concrete deliverables:

- *A set of tools that can interact and work together and demonstrate the achievable optimizations on a particular hardware platform.*

- *A methodology that enables the design of predictable embedded systems with a special focus on issues that cut several layers of abstraction, such as hardware and compiler design.*

Table of Contents

1. Overview of the Activity

1.1 ARTIST2 Participants: Expertise and Roles

Prof. Dr. Luca Benini - University of Bologna (Italy)
Power modelling and OS integration

Prof. Dr. Petru Eles - University Linköping (Sweden)
Dynamic and leakage power optimisation for real-time system, accurate system-level power modelling for communication

Prof. Dr. Tom Henzinger – EPFL (Switzerland)
Virtual machines for hard real-time computing, Embedded Machine; Schedule-Carrying Code

Prof. Dr. Rainer Leupers - RWTH Aachen (Germany)
Processor modelling tools

Prof. Dr. Jan Madsen, Technical University (TU) of Denmark (Denmark)
Power modelling and resource aware design space exploration

Prof. Dr. Peter Marwedel - University of Dortmund (Germany)
Memory architecture aware code generation

Prof. Dr. Reinhard Wilhelm - University of Saarbrücken (Germany)
Time as a resource.

Prof. Dr. Lothar Thiele – ETH Zürich (Switzerland)
Cooperating on concepts for real-time components (but not included in budget)

1.2 Affiliated Participants: Expertise and Roles

Roberto Zafalon – STM (Italy)
Dynamically controlling power consumption in MPSoC platforms.

1.3 Starting Date and Expected Ending Date

Starting date: September 1st, 2004.

Expected ending date (within ARTIST2): End of the project

Limited resources are a key characteristic of embedded systems. For example, no solution to the problem of very limited energy supply is in sight. Therefore, resource aware design will continue to play a key role well beyond the end of ARTIST2 funding.

1.4 Baseline

The importance of resource awareness in embedded systems is growing rapidly. The limited availability of computing resources is preventing the introduction of new products and applications, especially is areas where high-performance embedded systems are required (e.g. in telecom and consumer markets). Resources include energy, computational power and hardware components.

Minimization of the energy consumption plays a major role in the design of embedded systems. Limited availability of energy is the dominating constraint for many advanced embedded

systems, in particular those involving multimedia or sensor technologies. At the start of the project, only a very limited number of energy models were available. Optimizing systems for minimized energy consumption at the software level was more an exception. Exploitation of memory hierarchies for minimizing the energy consumption was at an infant state.

Also, maximizing the computational power of embedded systems is becoming of increasing importance due to the wide-spread deployment of high-performance multimedia-enabled devices in the market. The increasing trend towards encrypted communication also increases the performance requirements. As a result the efficient use of available hardware components, most importantly of processors and memories, is mandatory. When the project was started, customized processors were in use, but the generation of tools chains for processors was a mainly manual process. Support for multiprocessor systems was almost completely lacking.

With the growing software content in embedded systems and the diffusion of highly programmable and re-configurable platforms, software is given an unprecedented degree of control on resource utilization. This relation between hardware and software layers can be used to perform aggressive optimizations that can be achieved only by a synergistic approach that combines the advantages of static and dynamic techniques.

1.5   Problems Tackled in Year2

In general, a tight integration between the software and hardware layers is a key approach for coping with resource constrains. Research in year 2 addressed all embedded system resources which are available only in tightly limited amounts:

1. Coping with **energy constraints** has been a research goal in several groups involved in this activity and it has also been the goal of co-operations.

    The particular problem addressed by the Linköping group during this period has been that of on-line approaches for voltage scaling and dynamic body biasing. The goal was to address both dynamic and leakage power and, as opposed to static approaches, to also make use of the dynamic slack resulting from the fact that tasks execute for less than their worst case execution times (WCET). Shut-down of processors, very important for leakage power reduction, has also been considered. One of the main challenges was that such an on-line approach has to be of very low complexity.

    Dortmund improved the software support for memory hierarchies. Possible solutions include architectures containing small, fast memories called scratch pads and compilers that map memory objects to these memories at compile-time, rather than at run-time. Such compilers must be aware of the resource "memory hierarchy". In year 2, improvements for the existing optimization tools were the target. These improvements concerned the support for multiple processes, the integration of various models and their simulation and the generation of run-time support for scratch pad management.

2. Limited **computational power** is a second constraint.

    Highly optimized processors can be designed with tools that support the creation of tool chains for application-specific instruction set processors. This particular problem is tackled by the Aachen group by providing a method for generating tool chains for application-specific processors. In year 2, the problem tackled was to tighten the interface between the LISATek tools designed at Aachen and tools designed at other partners, in particular at Bologna. The goal included the accurate prediction of the performance of a complex multiprocessor system, enabled by the MPARM virtual platform from Bologna. This required the generation of models for hardware components typically found in systems on a chip. Towards this end, compatible models of interrupt and peripheral components had to be generated. Compatibility between the LISATek and MPARM environments had to be provided. Furthermore, efficient use of

memories is becoming more and more important as the speed gap between processors and memories widens. The DTU group has focused on resource aware multiprocessor design space exploration, taking into account memory, buffer and power constraints.

3. A third goal of the resource aware design activity is to provide **timing predictable** designs. Current computer architectures are optimized towards excellent average performance. For hard real-time systems, the worst case behaviour is what counts. As a consequence, for example, memory hierarchies improving the worst case behaviour and not just the average case behaviour are necessary. Timing predictability was tackled by six groups. The Linköping group focussed on WCET calculations for multiprocessors, while the two groups at Saarbrücken worked on generating tighter bounds for access times to caches. Support for scratch pads, as studied by Dortmund, also establishes tight bounds on the WCET. Analysing of these bounds has been a target in the second year. Improved knowledge about the real-time behaviour of components was the target of the work at ETH Zurich and EPFL. The six groups discussed ideas for designing predictable systems. A joint research proposal was filed by five of the six groups.

2. Summary of Activity Progress

2.1 Previous Work

Work achieved in the first 6 months:

Cooperation was established between the Universities of Bologna and Dortmund. The objective was to integrate the memory-aware compiler developed in Dortmund with the multi-processor platform simulator developed in Bologna. The first results were the definition of a standard format for the executable output of the compiler, as well as for the memory allocation information. This output format is supported by the platform simulator.

Cooperation was furthermore established between the Universities of Bologna and Aachen. The objective is to extend the modelling capabilities of the platform simulator developed in Bologna toward heterogeneous multi-core architectures, exploiting the Application-specific Processor development framework based on the LISA architecture description language developed in Aachen. The first result of this work was the definition of a standardized wrapping protocol which allows any SystemC core description generated by Aachen tools to be instantiated (multiple times) as a core in the platform simulator by Bologna.

Work achieved in months 6-12

The cooperations established in the first six months were continued and were significantly strengthened, as a significant amount of technical work was performed to sustain them. More specifically:

1. The cooperation between Bologna and Dortmund required the development of a new souce-level transformation tool for performing memory optimizations by Dortmund, and the development of compatible memory organization models by Bologna (including I and D caches as well as scratchpad memories).

2. The cooperation between Aachen and Bologna required extensive re-design of Bologna's core interfacing protocol within the platform simulator. On the other hand, Aachen has provided extensive technical support on Lisatek core wrapping architectures and toolsets.

An additional cooperation between Bologna and Saarland University was established. The objective of this cooperation is the exploitation of the platform simulator developed at Bologna, more specifically of the timing accurate core models incorporated in the simulator, as targets for the worst case execution analysis framework developed in Saarland University.

Milestones:

1. Established a working path between the memory-optimizing compiler developed at Dortmund and the platform simulator developed at Bologna. This tool interoperability path was validated during a visit of Dortmund's research staff to University of Bologna in July 2005

2. Development and extensive benchmarking of the LISATek-MPARM bridge: Several heterogeneous platforms were instantiated and tested. Performance analysis was performed.

2.2   Current Results

2.2.1 Technical Achievements / Outcomes / Difficulties encountered

Energy efficient time constrained systems
Power models as well as a simulation environment for validation have resulted from cooperation of the University of Linköping with the Bologna group. As the first step, an approach for mono-processor systems has been elaborated, implemented and published [And05].

During the last six months the efforts have concentrated on an extension of this approach to multiprocessor systems. This work is currently performed as part of the ARTIST mobility action in cooperation with the Dortmund group and will be continued into the following period.

Predictability in Multiprocessor SoC architectures
Besides being energy efficient and having a high performance, for many applications it is required that multiprocessor SoC implementations are highly predictable with respect to their timing behaviour. This problem has been addressed by the Linköping group during this period. While this issue has been previously investigated in the context of mono-processor systems, available results are inapplicable to modern multiprocessor architectures in which, for example, due to the shared memory access and shared buses, the individual WCETs of tasks depend on the global system schedule. Providing WCET guarantees and reliable schedules in this context is extremely challenging. It involves issues related to bus protocols and control, WCET analysis, system level scheduling and optimizations. With regard to the "classical" aspect of WCET analysis the group is building on the Symta/P tool from the Braunschweig group (a member of the ("execution platform" cluster). The Linköping group is also interacting with the Bologna group with regard to the issues of bus control.

This work is an effort started at the beginning of 2006. The overall concept has been elaborated, solutions have been developed and tools are under implementation. Publications and further results are expected in the following period.

Web site: http://www.ida.liu.se/~eslab/real-time.html

Integration of LISATek ISS models in SystemC and the MPARM virtual platform
The issues arising from the integration of LISATek ISS models in SystemC and the MPARM virtual platform have been investigated in more detail [Ang05], especially concerning the interaction with level one (L1) memories. A new MPARM functional model was developed to handle the L1 memory. It was also useful to cluster other functionality within the same block. The end result is called a "processor tile", comprising LISATek-generated SystemC model of the processor and the most tightly coupled components (see fig. 1).

The following component models were developed:

- a timer device,

- an emulated serial port,

- a simple interrupt controller.

The first component is vital if attempting to port an operating system. The second is very useful for debugging purposes; placing it next to IP cores, instead of in a shared location accessible to all system processors, has the advantage of allowing for independent input/output, and prevents debug traffic from spilling onto the system interconnect where it could pollute performance statistics. Finally, the interrupt controller is both a requirement of the other two devices and a crucial component to develop efficient synchronization mechanisms in multiprocessor systems. The controller is externally attached to a set of system-level wires which convey inter-core interrupts. On the IP core side, a simple interrupt handshaking protocol

was implemented at Bologna. In this protocol, the value of interrupt registers is copied on some LISATek core pins which are polled every cycle by the core to take proper action. The interrupt controller is memory mapped to let the core reset the pending interrupt flags and configure the masking status.
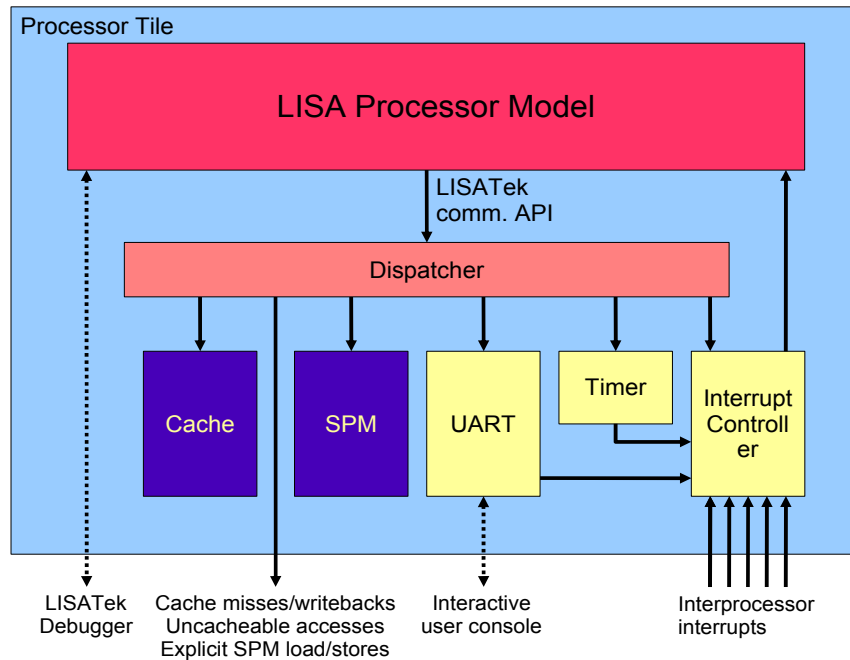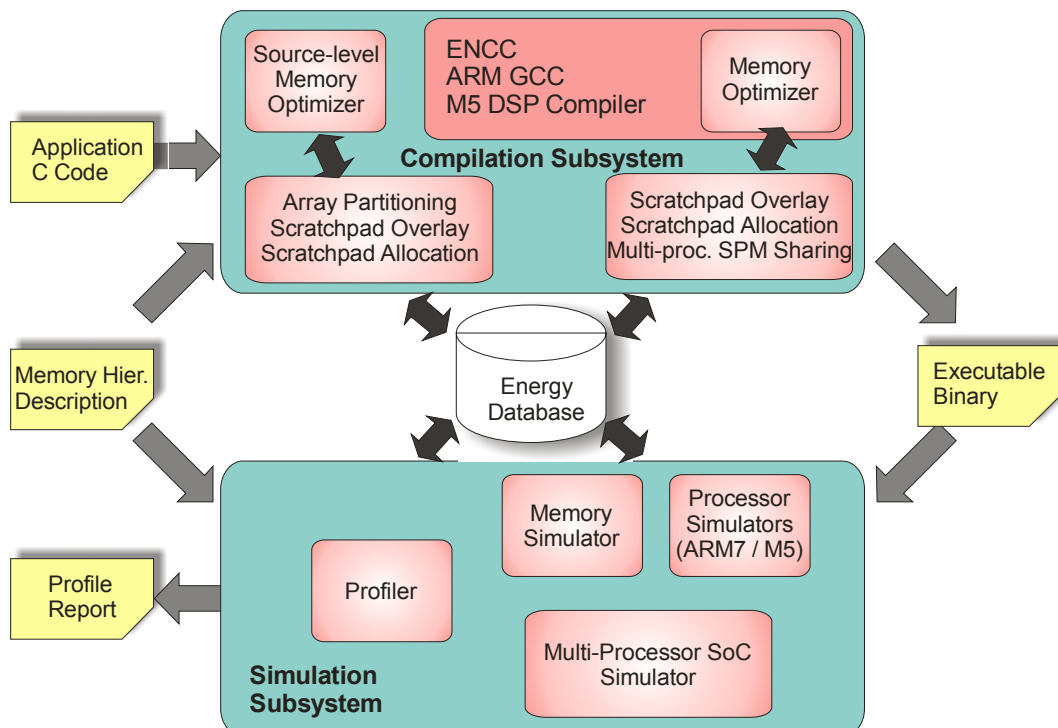

Figure 1 Processor Tile

Figure 2: Memory Aware Compilation and Simulation Tool-Chain

Memory Aware Compilation and Simulation Tool-Chain for Energy Optimizations

During the last reporting period, the need for a coherent tool chain for energy optimizations and for exploration of memory hierarchies across different system architectures was recognized. Therefore, a memory aware tool-chain supporting uni-processor ARM, multiprocessor ARM and M5 DSP based systems was developed (see fig. 2). Both the simulation and compilation subsystems are configured from a single memory hierarchy description. In addition, a common energy database is used by the memory optimizers in the compilation subsystem as well as by the memory and multi-processor SoC simulators in the simulation subsystem. The developed tool-chain optimizes input application code for a given memory hierarchy [Ver06d, Weh06] and also evaluates the optimization by simulating the optimized executable on the same memory hierarchy. The tool-chain is developed due to the cooperation between University of Dortmund and University of Bologna, as the simulation subsystem includes the multi-processor SoC simulation from Bologna while the compilation subsystem is developed at Dortmund. Moreover, both partners have agreed on a common memory hierarchy description format, which will be used for developing future optimizations [Ver06c].

Web site: http://ls12-www.cs.uni-dortmund.de/research/macc

Design-Time Memory Allocation Techniques for Multi-Process Applications with Aperiodic Processes

Previous work at Dortmund proposed compile-time or design-time memory allocation approaches to share the scratchpad memory among the periodic processes of a multi-process application. The current work extends the previous work and proposes memory allocation approaches for applications consisting of aperiodic tasks. This significantly increases the complexity of the memory allocator as the arrival times of the processes are completely unknown at design time. Therefore, the memory allocator is divided into an intelligent design-time component and a simple run-time component.

The design-time component of the memory allocator works in the following stepwise manner. First, it identifies memory objects, *i.e.* code segments and data variables, which on scratchpad allocation lead to reduction in the energy consumption of the system. Second, it processes the application code to enable the movement of memory objects at runtime. Finally, it inserts blocking statements in the application code to prevent unsafe movement of memory objects. The runtime component, depending upon the current set of active processes and the current state of the scheduled process, allocates (de-allocates) memory objects to (from) the scratchpad memory. Experiments report that a two-phased memory allocator minimizes the energy consumption due to applications with aperiodic tasks [Ver06a, Ver06b].

Web site: http://ls12-www.cs.uni-dortmund.de/research/macc

Operating System Support for Online Allocation of Scratchpad Memories

The goal of this work at Dortmund is to develop a runtime memory allocator which keeps track of the execution behaviour of the application and allocates scratchpad memory with memory objects (code segments and data variables) at runtime. The runtime allocator of this approach is more complex than the design-time memory allocator described above. At compile time, attributes such as access counts and the size are computed for each memory object. These attributes are then supplied as input to the memory allocator. The allocator based upon the input attributes, the scratchpad memory utilization and the current execution pattern swaps memory objects in and out of the scratchpad memory. Several heuristics as well as analytical approaches have been proposed for the online allocation of the scratchpad memory. The proposed approaches have been integrated into the RTEMS operating system. Experiments demonstrate that for highly dynamic applications, significant energy savings can be achieved.

Web site: http://ls12-www.cs.uni-dortmund.de/research/macc

Analysis of cache predictability

First quantitative results have been obtained on the predictability of different cache architectures. A paper is in preparation.

Improvement of timing analysis by integration with code synthesis

The University of Saarbrücken and AbsInt (an industrial member of the compiler cluster) have cooperated with ETAS (an external company) on the integration of the ASCET-SD model-based design tool with the AbsInt timing analyzer aiT. This work is continuing. A paper was published by Ferdinand et al. [Fer06].

Web site: http://en.etasgroup.com/about/tradeshows/documents/2006-03-15_AutomotiveSoftwareWorkshop_ASCET_Paper_Renz.pdf

Resource aware design space exploration

The Technical University of Denmark (DTU) has developed a multi-objective design space exploration environment based on the PISA environment for multi-objective optimization from the group of Lothar Thiele, ETH Zurich. The exploration is based on a genetic algorithm to solve the problem of mapping a set of task graphs onto a heterogeneous multiprocessor platform. The objective is to meet all real-time deadlines subject to minimizing system cost and power consumption, while staying within bounds on local memory sizes and interface buffer sizes. The approach allows for mapping onto a fixed platform or onto a flexible platform where architectural changes are explored during the mapping. This work will be continued. A paper has been accepted for publication at DIPES 2006 [Mad06]

FET Open Call project proposal

A consortium from within ARTIST2 consisting of the Universities of Saarbrücken, Zürich, Bologna, Pisa and Dortmund as well as AbsInt has applied for a project on "Reconciling Performance with Predictability" in the FET Open Call. Both short and long proposals have passed all thresholds. However, only 5% of the proposed projects can be funded, and this project will probably not be among them.

Interfaces for real-time components

Between members of the group of Tom Henzinger (EPFL) and Lothar Thiele (ETHZ) there have been intensive discussions on interface based design of embedded systems. There were common meetings and presentations. The main concept is to extend the common idea of static types towards resource types that talk about the use of various resources by a component, e.g. power, time, computing resources. As a result, the concept of interface-based design (by Tom Henzinger) has been successfully applied to real-time systems and associated publications have been written [Hen06, Thi06, Cha06].

Web site: http://chess.eecs.berkeley.edu/pubs/92.html

Resource awareness in sensor networks

The University of Bologna cooperated with ETH Zürich on resource awareness in sensor networks. For a full description please refer to the report on progress within the execution platform cluster.

Difficulty: LISATek Compatibility

LISATEK versions distributed by Europratice were frequently not based on recent versions of Linux. The cooperation between the Universities of Aachen and Saarbrücken would benefit from a shorter update cycle.

2.2.2 Publications Resulting from these Achievements

[And05] Alexandru Andrei, Marcus Schmitz, Petru Eles, Zebo Peng, Bahir M. Al-Hashimi: Overhead-Conscious Voltage Selection for Dynamic and Leakage Energy Reduction of Time-Constrained Systems, *IEE Proceedings Computers & Digital Techniques, Volume 152, Issue 1, 2005, pp. 28-38.*

[Ang06] Federico Angiolini, Jianjiang Ceng, Rainer Leupers, Federico Ferrari, Cesare Ferri, Luca Benini: An Integrated Open Framework for Heterogeneous MPSoC Design Space Exploration, *Proceedings of the Design, Automation and Test in Europe Conference and Exhibition 2006, Munich, Germany, Mar 6-10, 2006, pp. 1145-1150*

[Cha06] Samarjit Chakraborty, Lothar Thiele, Ernesto Wandeler, Nikolay Stoimenov: Interface-Based Rate Analysis of Embedded Systems, submitted to RTSS 2006.

[Fer06] Christian Ferdinand, Reinhold Heckmann, Hans-Joerg Wolff, Christian Renz, Oleg Parshin, Reinhard Wilhelm. Towards Model-Driven Development of Hard Real-Time Systems – Integrating ASCET-MD and aiT/StackAnalyzer. In *Proceedings of the Automotive Software Workshop in San Diego, 2006*, San Diego, USA, March 2006

[Hen06] Thomas A. Henzinger and Slobodan Matic: *An interface algebra for real-time components*, Proceedings of the 12th Annual Real-Time and Embedded Technology and Applications Symposium (RTAS), IEEE Computer Society Press, 2006.

[Mad06] Jan Madsen, Thomas K. Stidsen, Peter Kjærulf, Shankar Mahadevan: *Multi-Objective Design Space Exploration of Embedded System Platforms*. Accepted for publication: DIPES 2006, Portugal, 2006.

[Thi06] Lothar Thiele, Nikolay Stoimenov, Ernesto Wandeler: Real-Time Interfaces for Composing Real-Time Systems. To be published: EMSOFT 2006, Seoul, 2006.

[Ver06a] Manish Verma, Peter Marwedel: *Advanced Memory Optimization Techniques for Low-Power Embedded Processors.* In Fundamentals and Methods for Low-Power Information Processing (Ed. Baerbel Mertsching), Springer, Dordrecht, The Netherlands, 2006.

[Ver06b] Manish Verma: *Advanced Memory Optimization Techniques for Low-Power Embedded Processors,* PhD Thesis, University of Dortmund, Germany, 2006.

[Ver06c] Manish Verma, Lars Wehmeyer, Robert Pyka, Peter Marwedel, Luca Benini: *Compilation and Simulation Tool Chain for Memory Aware Energy Optimizations*, SAMOS 2006, 6th International Workshop, p. 279-288.

[Ver06d] Manish Verma, Peter Marwedel: *Overlay Techniques for Scratchpad Memories in Low Power Embedded Processors,* IEEE TVLSI, vol. 14, no. 8, August 2006

[Weh06] Lars Wehmeyer, Peter Marwedel: *Fast, Efficient and Predictable Memory Accesses*, Springer, 2006

2.2.3 Keynotes, Workshops, Tutorials

Keynotes: Peter Marwedel: Towards laying common grounds for embedded system design education, Opening, Embedded Systems Week (at Manukau Institute of Technology)
Auckland, New Zeeland, Nov. 16th, 2005
The talk proposed an approach for introducing embedded systems at the college level.

Mini-Keynote: Jan Madsen: Evolving MPSoC Solutions
MPSoC Symposium, Colorado

A key challenge of implementing an embedded systems application on a heterogeneous multiprocessor SoC platform is to find the right partitioning of the application onto the platform architecture. The right partitioning is dependent on the characteristics of the processors and the network connecting them, as well as the application. The mini-keynote addressed this challenge.

http://tima.imag.fr/MPSOC/.

Workshop: SCOPES: 9th International Workshop on Software and Compilers for Embedded Systems
Dallas, US – Sept. 29th – Oct. 1st, 2005
Software for embedded systems with emphasis on code generation for embedded processors.
http://www.scopesconf.org, http://www.edaa.com/activities.html

Tutorial: Peter Marwedel: Code optimizations for efficient embedded systems (at SCOPES 2005)
Dallas, US, Sept 29th, 2005
The tutorial presented various code transformations aiming at improving the efficiency of embedded software, taking the limited resources of embedded systems into account
http://www.scopesconf.org/Pgm.pdf

Tutorial: Luca Benini: System Level Power Optimization (at course on Advanced Digital Design, organized by EPFL)
Lausanne, Switzerland, Oct. 8th, 2005
The tutorial presented the main issues in power optimization (under various types of resource constraints) at the system level. The tutorial aimed at industrial as well as academic attendees.
http://www.mead.ch/CoursLAP/ADSD-Program.html (for the 2006 edition)

Tutorial: Rainer Leupers: Retargetable Compilation (at course on Advanced Digital Design, organized by EPFL)
Lausanne, Switzerland, Oct. 6th, 2005 (morning)
The tutorial presented techniques for generating compilers from descriptions of the instruction set architecture (ISA). The tutorial aimed at industrial as well as academic attendees.
http://www.mead.ch/CoursLAP/ADSD-Program.html (for the 2006 edition)

Tutorial: Peter Marwedel: Memory-architecture aware compilation (at course on Advanced Digital Design, organized by EPFL)
Lausanne, Switzerland, Oct. 6th, 2005 (afternoon)
The tutorial presented the benefits resulting from making compilers aware of the memory architecture. Significant reductions in terms of consumed resources (energy, time) can be achieved. The tutorial aimed at industrial as well as academic attendees.
http://www.mead.ch/CoursLAP/ADSD-Program.html (for the 2006 edition)

Tutorial: Peter Marwedel: Code optimizations for efficient embedded systems (at Manukau Institute of Technology)
Auckland, New Zealand, Nov. 17th, 2005
The tutorial presented various code transformations aiming at improving the efficiency of embedded software, taking the limited resources of embedded systems into account.

Tutorial: Lothar Thiele: Frameworks for System-Level Analysis of Real-Time Systems - Symta/S and MPA

RTAS 2006 Tutorial IEEE Real-Time and Embedded Technology and Applications Symposium: System-level timing, performance, and power becomes increasingly intractable as the interactions between system parts introduce complex dynamic behaviour that can not be fully overseen by anyone in a design team. The tutorial addressed recent research on composable and extensible analysis methods, and tools.

Tutorial: Lothar Thiele: Sensor Networks
DATE 2006 Symposium
This tutorial reviewed basic concepts of wireless sensor networks, including: ad-hoc networking, programming models, power management, in-network processing, development environments and methodologies.

Tutorial: Lothar Thiele and Peter Marwedel: ARTIST2 Spring School in China on Models, Methods and Tools for Embedded Systems
Xi'an, China, April 3rd-15th, 2006

The tutorial started with an introduction to embedded systems and resource aware generation of software and performance analysis. It also comprised modelling of real-time systems, validation and verification.
http://www.artist-embedded.org/FP6/ARTIST2Events/Events/ChinaSchool/

3. Future Work and Evolution

3.1 Problems to be tackled over the next 18 months (Sept 2006 – Feb 2008)

The Linköping group will continue to work in the context of predictability on multiprocessor SoC. An integrated environment will be implemented including WCET estimation, system scheduling and optimisation, taking into account the system wide interactions. Furthermore, the cooperation of the Linköping group with Braunschweig on WCET analysis and Bologna on the hardware implementation will continue. Cooperation with Dortmund on frequency and voltage scaling will continue.

Over the next reporting period, efficient utilization of memory hierarchies by multi-process applications will be the focus of the research at Dortmund. In the coming months, the group at Dortmund will enhance the memory allocation approaches to improve the scratchpad utilization by multi-process applications comprising of periodic and aperiodic tasks. A further improvement of the proposed approaches will be achieved by supporting applications with cooperative multi-tasking. The group also plans to develop memory allocation approaches as part of its compilation framework which is tightly integrated with the multi-processor SoC (MPARM) simulator from University of Bologna. The integration of the highly configurable memory hierarchy simulator developed at Dortmund into MPARM is also planned for the coming months. It is expected that these research efforts will further strengthen the cooperation between University of Dortmund and University of Bologna.

The group at the University of Saarbrücken will extend its estimation techniques for WCET bounds towards more complex architectures. Results will be integrated by partners.

The group at Bologna will continue to cooperate with STM on problems of industrial relevance. Furthermore, extended cooperation with Dortmund on memory modelling is on the horizon.

EPFL and ETH Zürich are planning to continue working on developing interface formalisms and algorithms for interface compatibility checking for interfaces that expose timing and resource constraints of components. Concretely, the partners hope to understand better the differences and commonalities between their interface formalisms, in order to combine or generalize them.

ETH Zürich will be working together with partners such as Bologna, Saarbrücken and Dortmund on an informal basis, in particular on the design of timing-predictable systems.

The Technical University of Denmark will extend its work on resource aware design space exploration. Furthermore, the cooperation with University of Bologna on system modelling infrastructure will be extended to include power models for more accurate exploration of power consumption. Finally, a proposal for a joint project between Technical University of Denmark and University of Linköping on analysis and synthesis of low-power fault-tolerant embedded systems has been filed. The aim is to consider efficient inclusion of fault tolerance given the tight resource constraints.

3.2 Current and Future Milestones

Year2: A set of tools that can interact and work together and demonstrate the achievable optimizations on a particular hardware platform.

Year 3: Use of the interacting tools, analysis of their potential and conclusions for future design methodologies; this includes, in particular, the use of the MPARM platform and integrated tools designed at the Universities of Aachen and Dortmund

Year4: A methodology for resource-aware design of embedded systems; this methodology will include mechanisms for handling optimizations across different levels of abstraction

3.3 Indicators for Integration

Integration indicators are extremely strong: several experiments demonstrate the effectiveness of resource usage optimization at various levels of the software development pipeline. This result has required considerable technical effort and demonstrates both the commitment of the partners and the potential for the cooperation. The following tools are interacting:

- Aachen / Bologna: LISATek / MPARM integration

- Aachen / Saarbrücken: Saarbrücken using LISATek tools.

- Bologna / Dortmund: MPARM simulation platform / memory-aware C-compiler macc, MPARM simulation platform / MEMSIM simulator.

  In particular, the cooperation between University of Dortmund and University of Bologna resulted in the development of a coherent memory aware compilation and simulation framework. The compilation framework includes several memory optimizations and can be obtained from the following URL: http://ls12.cs.uni-dortmund.de/ research/macc

  Both partners have also agreed on a common format to describe multi-processor SoC architectures. A common single format can be used to configure both the compilation and the simulation subsystems of the framework. Currently, the memory hierarchy simulator from Dortmund is being integrated into the multi-processor SoC simulator from Bologna.

- Bologna / Linköping: Using power models and simulation from Bologna within the Linköping environment

- Bologna / STM: energy estimations have been integrated into models from Bologna.

- Dortmund / Linköping: Voltage scaling implemented by a student from Dortmund at Linköping

- Dortmund / Saarbrücken (AbsInt): Tight integration described in the compiler cluster report.

- ETH Zürich / EPFL: Importing concepts for the description of real-time components and working towards an integration of the concepts.

- Linköping / Braunschweig (cluster member for execution platforms): building on top of Symta/P from Braunschweig at Linköping

- Linköping / TU Denmark: extending the ARTS framework from TU Denmark at Linköping towards a framework for evaluation of the degree of pessimism introduced by WCET analysi, in particular in the context of distributed embedded systems as found in automotive applications.

- TU Denmark / Bologna: using power models and simulation from Bologna at TU Denmark.

3.4 Main Funding

Bologna: STMicroelectronics, direct industrial grant

Bologna: Freescale semiconductors, direct industrial grant

Swiss National Science Foundation (Estimation of System Properties)

Deutsche Forschungsgemeinschaft (DFG)

Swedish Foundation for Strategic Research (SSF)

3.5 Internal reviewers

- Dr. Lars Wehmeyer, Informatik Centrum Dortmund e.V. (ICD), Dortmund

- Prof. Dr. Gernot Fink, Informatik 12, Perception group, Universität Dortmund