IST-004527 ARTIST2
Network of Excellence
on Embedded Systems Design

Cluster Progress Report for Year 3

Cluster:

Execution Platforms

Cluster Leader:

Prof. Dr. Lothar Thiele (ETH Zurich)
http://www.tik.ee.ethz.ch

Policy Objective (abstract)

This topic is strongly linked to the compilation and implementation of embedded systems. For a given application, it is important to have the technology, methods and tools to make rational choices about the platform and the design used, before proceeding to final implementation. Research in Execution Platforms targets the development of the theoretical and practical tools for modelling the dynamic behaviour of application software for a given platform. This is a new area of research, which will allow greater flexibility in designing optimal embedded systems.

Table of Contents

1. Overview 3
1.1 High-Level Objectives 3
1.1.1 System Modelling Infrastructure 3
1.1.2 Communication-Centric Systems 4
1.1.3 Low-Power Design 4
1.1.4 Resource-Aware Design 5
1.2 Industrial Sectors 6
1.2.1 Automotive Industry 6
1.2.2 Mechatronic Industry 6
1.2.3 Information Technology 7
1.3 Main Research Trends 7

2. State of the Integration in Europe 10
2.1 Brief State of the Art 10
2.1.1 Communication-Centric Systems 10
2.1.2 Resource-Aware Design 12
2.1.3 Simulation and Performance Analysis 13
2.2 Main Aims for Integration and Building Excellence through Artist2 14
2.3 Other Research Teams 14
2.4 Interaction of the Cluster with Other Communities 16

3. Overall Assessment and Vision for the Cluster 18
3.1 Assessment 18
3.1.1 Integration Activities within the Cluster 18
3.1.2 List of Joint Publications 19
3.2 Vision and Long Term Goals 23
3.3 Plans for Year 4 23

4. Cluster Participants 26
4.1 Core Partners 26
4.2 Affiliated Industrial Partners 30
4.3 Affiliated Academic Partners 31
4.4 Affiliated International Partners 32

5. Internal Reviewers for this Deliverable 33

1. Overview

In relation to the ARTIST2 network, it is the overall goal of the topic on execution platforms to extend the current state in composability towards issues like modeling of non-functional constraints, power and energy, end-to-end real-time behavior, timing and performance analysis and heterogeneous models of computation.

One of the most critical issues to be faced in the research on execution platform is their rapidly growing complexity. Complexity increase is pushed by Moore's law, and by the ever-increasing demand for high performance computing. In addition, the boundaries between hardware and software domains are getting more blurred (dynamically re-configurable hardware, adaptable embedded systems, breakthrough in power consumption and performance) and challenging questions are at the border (trade-offs in mapping applications to hardware and software components, re-configurable hardware, WCET, performance of distributed computer and communication systems).

The cluster attempts to combine the diverse knowledge and to integrate different approaches approaches in the area of execution platforms for embedded systems available in Europe and beyond.

1.1 High-Level Objectives

The workplan of the platform can be partitioned into System Modelling Infrastructure, Communication-Centric Systems, Low-Power Design and Resource-Aware Design.

1.1.1 System Modelling Infrastructure

The aim is to provide a scalable and realistic modelling platform which is abstract enough to provide complete system representations and some form of functional models even for billion-transistor future systems, while at the same time providing the needed flexibility for modelling a number of different embodiments (e.g. multi-processors, homogeneous and heterogeneous, reconfigurable, etc.).

In year 1 the focus was on assessing the state-of-the-art within the two different modelling approaches, simulation- and formal-based. In year 2 the integration within simulation- and formal-based modelling were further researched and in particular modelling extensions which allow combinations of simulation- and formal-based modelling approaches were investigated. The focus for year 3 (covered in the “next 18 months” section from year 2) was to extend the simulation-based modeling to address issues of dynamically reconfigurable architectures, distributed embedded systems, and lab-on-a-chip, and to consolidate the the integration between the simulation models of Bolgna and DTU. For the formal-based modeling the focus was on model integration, using timed automata models, and to extend current models to address and encompas more hardware issues.

These objectives have been achieved. For the simulation-based modeling approaches, the ARTS model from DTU has been extended with capabilities for simulating the dynamic behavior of run-time reconfigurable platforms. LiU have finalized their distributed embedded systems simulator derived from the ARTS model, and have used it to validate formal analysis approaches of CAN and FlexRay based systems, by estimating their degree of pessimism. Based on their MPARM simulation framework, Bologna (together with EPFL) have developed a NoC emulation framework which acts as a design tool for tuning and functional validation of on-chip interconnections for MPSoCs. This emulation framework is implemented onto a *Field Programmable Gate Array* (FPGA) platform in order to increase analysis performance. Finally, the model integration of MPARM and ARTS, and the jointly developed reactive traffic generators have been finalized. Within the formal-based modeling approaches, Braunsweigh has extended their timing analysis to address shared resources such as memories and busses, issues which are important for modeling MPSoC platforms. Braunsweigh have extended their sensitivity analysis to handle multi-dimensional analysis and have proposed expressive robustness metrics for different assumptions and design scenarios, and showed how these can be efficiently considered throughout the whole design process. ETHZ have focused on integrating their MPA modeling formalism with other system models. In particular there has been a deep integration between Symta/S and MPA, which not only entails converters between the two modeling formalisms but also investigations, when to use which formalism. Finally, DTU have expressed most parts of their ARTS model in the timed automata semantics of UPPAAL and have demonstrated the verification of applications executing on a multi-core execution platform.

Based on the reviewers recommendation to address issues of fault-tolerance of embedded systems, LiU and DTU have proposed an extension to the process graph model which is able to capture the occurrence of faults. They have shown how design transformations that introduce redundancy, such as re-execution and replication, can be applied on this model. LiU have extended their modeling of NoC platforms to address transient faults, allowing them to optimize NoC communication for fault-tolerance.

1.1.2 Communication-Centric Systems

The work aims at new best-case/worst-case models for hard real-time systems and at combined statistical and interval models for QoS applications in multi-media. These models will combine communication and computation, different models of computation, event models and scheduling policies.

In the first 12 month the state-of-the-art in models was assessed taking into consideration the particularities of various (quasi)standard communication protocols during system analysis and scheduling. In the second 18 months the modelling scope was extended to cover emerging hierarchical protocols such as FlexRay in the automotive domain, and to include the new aspect of fault tolerant systems.

In the first 12 month of the project, a report was produced unifying and relating the different approaches to communication-centric systems and low power design. In the second 18 month we coupled several tools to extend the analysis scope and accuracy. First, the power models for scheduling were extended, second the simulation platforms for communication centric systems, MPARM and ARTS, and third, MPA and SymTA/S, were integrated. Additionally, we started new activities towards fault tolerance in heterogeneous embedded systems, including multi-core systems.

The objectives have been met in joint efforts by the respective teams. In the next period, the work in fault tolerant systems will be intensified in a collaboration of DTU and LiU. Fault tolerance activities will be further strengthened in a project that develops strategies for autonomous multi-core architectures. That project also includes several research partners external to the project. Timed automata will be considered as a further modelling construct to be coupled to MPA and SymTA/S.

1.1.3 Low-Power Design

The high-level objectives are in two directions: strengthening integration and making inroads in effective techniques for system-level power optimization.

Strengthening integration: work will be performed aiming at defining abstractions and models suitable for very-high-level system power estimation, both for localized and distributed platforms. Example platforms will be used to drive the integration effort, in the domain of multi-processor-systems on chip (leveraging the MPARM infrastructure) and in the domain of distributed sensor networks (leveraging the ARTS simulation model and the prototype sensor networks).

System-level power optimization: in this area, the interaction between resource allocation and scheduling and power optimization will be explored both for localized and distributed systems. Furthermore, the interplay between power and other cost metrics (reliability, performance) will be explored.

These objectives are actively pusued, and several milestones have been achieved leveraging the cooperation of the groups involved in the activity. More specifically: simulation-based power estimation, based on the MPARM infrastructure has been considerabily strengthened. Power models for interconnect fabrics and on-chip memories have been developed and extended to account for the characteristics of deep submicron technologies (leakage effects, variable voltage supply support). From the point of view of power-aware allocation and scheduling, significant progress has been made both in considering complex and realistic platform models (multi-core architectures, with variable frequency, shutdown support), and in the development of aggressive ostrategies to find optimal allocation and mappings of task-based parallel applications. Recently, also temperature-related effects have been taken into conisderation. The approaches have been extended to handle not only hard but also soft real-time systems. Innovative approaches to energy optimisation for wireless sensor networks have been developed.

1.1.4 Resource-Aware Design

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

> 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 is to be found.

Significant progress toward reaching these objectives has been made, especially on the front of the integration between research infrastructures. Important milestones achieved in this area are: (i) the construction of a common simulation-based infrastructure for modelling, analysis and design space exploration of heterogeneous multi-core architecture. The integration work by UNIBO and AACHEN has continued in the third year, focusing on advanced processor architectures capable of issuing multiple instructions in a single cycle. In this area a new milestone has been reached, namely the development of a complete very-long instruction word model in LisaTEK, and the inclusion of its simulation model within the MPARM platform. Moreover, a synthesizable model has been derived and fully synthesized in CMOS 65 techology. (ii) the integration of a optimizing compiler infrastructure (developed in UDORTMUND) with the virtual platform developed in UNIBO. Work has continued in this area, in the direction of creating an MPARM compiler back-end to support code parallelization. A milestone has been reached in this area: namely, complete support for barrier-based parallelization has been implemented and demonstrated.

1.2 Industrial Sectors

1.2.1 Automotive Industry

The automotive industry is currently in a fast and spectacular evolution towards the intelligent, safe, environmental, interconnected, and economic car. Electronics is at the basis of most of this development. New features such as automatic intelligent parking assist, blind-spot information system, navigation computers with real-time traffic updates, car-to-car communication, not to mention electronically controlled brakes or electronic power steering, are out and running in most recent high-end cars. This development is going to continue with new functionality being adopted not only in premium cars but also in the mass-market. Consequently, estimates are that up to 80% of the innovations are directly dependent on embedded systems. This, of course, comes at a cost. Automotive electronics, currently accounts for 22% of a vehicle's cost and is predicted to increase to 40% by 2010 (www.altera.com).

This evolution brings a series of challenges in all steps of the development cycle. How to specify and model such a complex system? There is a need for a component based modelling, analysis, and synthesis approach in which independently designed hardware and software components can safely be combined into a working system. How to achieve the ever increasing demand on functionality and safety, at an affordable cost? Modern automotive electronic systems are highly distributed networks with components interacting over various infrastructures. How to achieve a safe and predictable system at such a huge level of complexity and heterogeneity? A well defined methodology is needed for mapping the complex functionalities on predefined distributed automotive platforms. This assumes well defined standards, middleware layers, analysis tools, software generation tools, design exploration and optimisation approaches.

Due to ARTIST2 activities, (e.g. the ARTIST workshop “Beyon AUTOSAR” in Innsbruck) several technical meetings between TU Braunschweig and leading automotive suppliers in the AUTOSAR context held place. As a main topic it was discussed how compositional performance verification methods can be utilized in the automotive design process to facilitate the network integration problem. TU Braunschweig was invited to the SAE world congress 2007 in Detroit to present recent results in compositional performance verification. In addition, in cooperation with Daimler, the ESI actively participated in defining the Artemis research agenda in this area.

1.2.2 Mechatronic Industry

Traditionally, the development of mechatronic systems was a rather sequential process. First the mechanical part was designed, next the hardware infrastructure was fixed, and finally the embedded software was developed. Typically, this leads to many proplems at systems engineering, because only then the interference of design decisions from the disciples became visible. To improve this process and to shorten the time-to-market, there is a clear trend towards concurrent engineering. To be able to detect problems earlier in the development cycle, there a strong need for high-level models allowing early analysis of system-level design decisions. Moreover, there is an increasing interest in the use of models to improve the early testing process; for instance, one would like to test the embedded software before its environment is available.

Concerning the execution platforms used, one can observe the need for a flexible process where one can easily switch between various solutions, such as the amount of distribution, the topology used, the communication infrastructure, and the operating system. Often in a first release of a high-tech system the execution platform is overdimensioned. For instance, one might choose a highly distributed architecture to avoid scheduling problems. In a later version, a strong cost reduction has to be achieved by combining more functionality on a single node. One major problem is to foresee at an early stage of the design whether a particular hardware platform is feasible for a given software system. Hence there is a strong need for methods that can help engineers to make a well-founded choice for an execution platform.

An increasing interest in the application of model-driven design techniques can be observed. These techniques emphasize the explicit separation between the application logic and the execution platform and allow models to be analyzed and systemically refined through model transformations.

1.2.3 Information Technology

Microelectronic technology is continuing to grow according to Moore's law. However, the need for computation power in industry is growing even faster. This is the case with traditional areas such as technical/scientific computation, and, more recently, modern applications, for instance interactive multimedia, high bandwidth communication, or speech recognition. Many of these applications are running on mobile computer, which makes issues even more complicated: an unprecedented amount of computation power has to be delivered with very low energy consumption. So, instead of just running after high performance, industry is out after a good performance - energy product. These unprecedented performance/energy requirements cannot be achieved by further pushing processor technology along the traditional Pentium lines. New architectures are needed in which several lower performance (and less energy hungry) computation nodes are cooperating in order to globally achieve the expected performance. Modern MPSoC and NoC architectures are developed along these lines.

Another clear trend is towards reconfigurable architectures, in general, and configurable processors, in particular. The generic goal is to achieve a high degree of flexibility (traditionally available only with software implementation) at a power consumption which is much lower than achievable with a traditional software implementation using general purpose processors.

The emerging trend for multimedia applications on mobile terminals, combined with a decreasing time-to-market and a multitude of standards have created the need for flexible and scalable computing platforms that are capable of providing considerable (application specific) computational performance at a low cost and a low energy budget.

Hence, in recent years, the first multiprocessor System-on-Chip components have emerged (like e.g. TI OMAP, ST Nomadik, Philips Nexperia, IBM/Toshiba/Sonys CELL). These platforms contain multiple heterogeneous, flexible processing elements, a memory hierarchy and I/O components. All these components are linked to each other by a flexible on-chip interconnect structure. These architectures meet the performance needs of multimedia applications, while limiting the power consumption.

To effectively utilize these emerging technologies, new design methodologies are being developed. This includes application and architecture modelling, mapping but especially also design-space exploration techniques that aid in finding optimal trade-offs.

1.3 Main Research Trends

Many embedded system applications are implemented today using distributed architectures, consisting of several hardware nodes interconnected in a network. Each hardware node can consist of a processor, memory, interfaces to I/O and to the network. The networks are using specialized communication protocols, depending on the application area. For example, in the automotive electronics area communication protocols such as CAN, FlexRay and TTP are used. One important trend today is toward the integration of multiple cores on the same chip, hence embedded systems are not only distributed across multiple boards or chips, but also within the same chip.

As the complexity of the functionality increases, the way it is distributed has changed. If we take as an example the automotive applications, initially, each function was running on a dedicated hardware node, allowing the system integrators to purchase nodes implementing required functions from different vendors, and to integrate them into their system. Currently, number of such nodes has reached more than 100 in a high-end car, which can lead to large cost and performance penalties. Moreover, with the advent of poly-core (i.e. high cardinality multicore) single-chip platforms, the effective number of processing nodes tends to grow in a “fractal” way, and future distributed systems with thousands of processing nodes are not a far away dream.

Not only the number of nodes has increased, but the resulting solutions based on dedicated hardware nodes do not use the available resources efficiently in order to reduce costs. For example, it should be possible to move functionality from one node to another node where there are enough resources (e.g., memory) available. Moreover, emerging functionality, such as brake-by-wire, is inherently distributed, and achieving an efficient fault-tolerant implementation is very difficult in the current setting.

Moreover, as the communications become a critical component, new protocols are needed that can cope with the high bandwidth and predictability required. The trend is towards hybrid communication protocols, such as the FlexRay protocol, which allows the sharing of the bus by event-driven and time-driven messages. Time-triggered protocols have the advantage of simplicity and predictability, while event-triggered protocols are flexible and have low cost. A hybrid communication protocol like FlexRay offers some of the advantages of both worlds. The need for scalable and predictable communication is not only a characteristics of automotive designs, but even multimedia and signal processing systems are increasingly communication dominated.

With growing embedded system complexity more and more parts of a system are reused or supplied, often from external sources. These parts range from single hardware components or software processes to hardware-software (HW-SW) subsystems. They must cooperate and share resources with newly developed parts such that the design constraints are met. There are many software interface standards such as CORBA, COM or DCOM, to name just a few examples that are specifically designed for that task. Nevertheless in practice, software integration is not a solved but a growing problem. This is especially true when performance and energy efficiency can be achieved only if a sufficient degree of parallelisim in application execution is achieved.

New design optimization tools are needed to handle the increasing complexity of such systems, and their competing requirements in terms of performance, reliability, low power consumption, cost, time-to-market, etc. As the complexity of the systems continues to increase, the development time lengthens dramatically, and the manufacturing costs become prohibitively high. To cope with this complexity, it is necessary to reuse as much as possible at all levels of the design process, and to work at higher and higher abstraction levels, not only for specification of overall system functionality, but also for supporting communication among a number of parallel executing nodes.

One of the most significant achievements in the cultural landscape of low-power embedded systems design is the consensus on the strategic role of power management technology. It is now widely acknowledged that resource usage in embedded system platforms depends on application workload characteristics, desired quality of service and environmental conditions. System workload is highly non-stationary due to the heterogeneous nature of information content. Quality of service depends on user requirements, which may change over time. In addition, both can be affected by environmental conditions such as network congestion and wireless link quality.

Power management is viewed as a strategic technology both for integrated and distributed embedded systems. In the first area, the trend is toward supporting power management in multi-core architectures, with a large number of power-manageable resourcers. Silicon technology is rapidly evolving to provide an increased level of control of-on chip power resources. Technologies such as multiple power distribution regions, multiple power-gating circuits for partial shutdown, multiple variable-voltage supply circuits are now commonplace. The challenge now is how to allocate and distribute workload in an energy efficient fashion over multiple cores executing in parallel. Also, one open issue is how to cope with the increasing amount of leakage in nanometer technologies, which tends to over-emphasize the cost of inactive logic, unless it can be set in a low-power idle state (which in many cases implies storage losses and high wakeup cost).

In the area of distributed low-power systems, wireless sensor networks are the key technology drivers, given their tightly power constrained nature. One important trend in this area is toward “battery free” operation. This can be achieved through energy storage devices (e.g. super-capacitors) coupled with additional devices capable of harversting energy from environmental sources (e.g. solar energy, vibrational energy). Battery-free operation requires carefully balancing harvested energy and stored energy agains the energy consumed by the system, in a compromise between quality of service and sustainable lifetime.

2. State of the Integration in Europe

2.1 Brief State of the Art

2.1.1 Communication-Centric Systems

With growing embedded system complexity more and more parts of a system are reused or supplied, often from external sources. These parts range from single hardware components or software processes to hardware-software (HW-SW) subsystems. They must cooperate and share resources with newly developed parts such that the design constraints are met. There are many software interface standards such as CORBA, COM or DCOM, to name just a few examples that are specifically designed for that task. Nevertheless in practice, software integration is not a solved but a growing problem. In a recent interview with the German weekly Magazine “Der Spiegel”, Jürgen Schrempp (DER SPIEGEL 14/2005, 02.04.2005), CEO DaimlerChrysler, has explained that systems integration was a key problem in the recent series of call backs of Mercedes cars.

Software and communication layers together with interface standards increase software portability but they do not yet solve some of the key embedded systems challenges, (1) integration verification and (2) the control of performance and other not functional constraints, such as power consumption or dependability. Integration verification is a general systems design problem that deserves much attention but is not the focus of the activity. Performance problems originate in the fact that platform components have limited performance and memory resources. Sharing resources introduces additional, not functional dependencies. Such dependencies slow down communication or processing and increase event jitter potentially leading to buffer over- or underflows and lost messages or system failures. Such failures are difficult to identify, because many of them appear as timing anomalies that are not discovered using the typical component test patterns.

There are several approaches to cope with large systems integration. Frequently, designers use simulation or test and add some design rules of thumb for uncovered cases. An example is to limit the load on a prioritized bus to 30 or 60% of the maximum bus load. This rule has a realistic formal background and can be backed by a proof by Liu and Layland for periodic systems that hold under certain circumstances on a single processor or bus. However, these approaches do not apply to multi-hop networked systems.

A more general approach is a conservative design style that decouples the integrated components and subsystems by assigning them a fixed share of the resources. This sharing uses a TDMA technique (Time Division Multiple Access) where fixed time slots are assigned to processes or communication. Unlike round-robin scheduling, unused time slots stay empty avoiding possible buffer overflow due to extra resource shares that speed up processing or communication. In essence, TDMA reaches a complete decoupling at the cost of sub optimal resource, and possibly energy, utilization. The TDMA technique can be extended all the way to software development, where the elegantly simple mathematical formulation describing TDMA performance can be used for a system wide performance analysis and control, such as in the Giotto tool of UC Berkeley.

Such a conservative design style is typical for the aircraft industry which uses bus protocols such as TTP. TTP has additional synchronization functions that support fault detection and fault tolerance. Recently, TTP has also been introduced to the automotive mass market where it shall replace the fixed priority CAN bus in the mid term. Audi is considering TTP as their future automotive bus standard while the majority of automotive companies have committed to

FlexRay which uses a TDMA protocol as a basis and runs a dynamic “mini slot” protocol in one of the time slots. This hierarchy offers the rigorous conservative design style for part of the automotive functions and a dynamic schedule with higher resource utilization but complex non functional dependencies for other functions that are less time critical. The practical impact of the relatively complex FlexRay protocol will be seen soon when the first automotive system developments will use this bus. But even when commercial tools for FlexRay will be available in the future, they will fall short in evaluating general multi-hop networks as they appear today, from automotive to MpSoC.

Another approach is formal timing analysis. The scheduling effects of fixed priority scheduling, such as used for CAN buses, are well understood and formal analysis has found its way into commercial tools, such as the Volcano software and tool set. For a while, these techniques were sufficient for fixed priority scheduling approaches used in many of the bus protocols and real-time operating systems. With the advent of more complex networks including bridges, “local” analysis techniques that look at a single bus or processor only are not sufficient any more. This development can be observed from distributed real-time systems down to the level of an individual chip where several different buses are combined today. On the chip-level, the problem is even more complicated due to the introduction of caches (see roadmap). At the same time, the individual bus protocol becomes more complicated as visible in the transition from the CAN to the FlexRay protocol. Hence, embedded real-time system design is approaching a new level of complexity for which the current tools and formalisms are not appropriate any more.

There are proposals combining few different scheduling strategies, e.g. RMS on a processor and TDMA on the bus. These are called holistic approaches and were introduced by Tindell. A very good recent example that shows the power of this „holistic“ approach is the work in ARTIST NoE by the group of Eles. This work covers many different scheduling techniques and can be considered as defining the state of the art in holistic analysis. Other work by Palencia and Harbor, again in ARTIST, is leading in holistic analysis of systems with task dependencies. In general, it appears more efficient to identify solutions that encompass the whole system than to consider local scheduling individually.

On the other hand, there is an apparent scalability problem when considering the huge number of potential subsystem combinations that require adapted holistic scheduling. An alternative is a modular analysis technique that combines local analysis of individual components using event models as interfaces. Local analysis and event model propagation are iterated for a global analysis. ARTIST is the center of gravity for the development of this composition technique with the two cooperating groups from ETH Zürich and TU Braunschweig working on different versions of this approach.

Better modeling, analysis, and optimization are only part of the solution to master communication-centric systems. Networked systems based on the extension of conventional communication networks suffer from increasingly complex and decreasingly predictable real-time behavior. Therefore, the second focus of this ARTIST 2 activity is the development of new predictable and scalable networks with an emphasis on networks-on-chip. This is the interest of DTU and University of Bologna.

Recently, there has been a substantial interest in the area of massively distributed embedded systems that are used for communication, sensing and actuating. These sensor networks appear in many different application domains such as environmental monitoring, disaster management, distributed large scale control, building automation, elderly care and logistics. This quickly emerging research area is closely related to the subjects of the cluster on execution platforms. In particular, we are faced with resource constraints in terms of power and energy, computing, memory and communication. In addition, some of the potential application areas are characterized by (hard) real-time requirements. The necessary design and analysis methods clearly ask for an integrated view of the whole distributed systems, i.e. taking into account hardware and software, a cross-layer view on the different protocols and algorithms, and new concepts for application specification, middleware and operating systems. It is not possible to describe in a few sentences the major challenges in this new kind of systems. Please find below an incomplete set specific to the area of execution platforms:

- Deployment of large distributed networks with limited communication and computation resources.
- Low energy computation and communication.
- Fault tolerance and reliability.

Currently, it is not clear whether this kind of new embedded systems is finding wide acceptance in industry. Nevertheless, there are already first projects in Europe that attempt to do a technology transfer from academia to industry, e.g. a safety network for building applications (Siemens, ETH Zurich).

2.1.2 Resource-Aware Design

While microelectronic technology is continuing to provide growth according to Moore's low, our need for computation power is growing even faster. Therefore, although, they have unprecedented computation and memory resources at their disposal, designers of embedded systems have to be increasingly aware of the fact that, in order to achieve the requirements of novel applications, these resources have to be used in the most efficient way.

A particular dimension of this problem is that of power consumption, since energy is one of the dominating constraints, especially, but not only, in the case of mobile applications. Reducing energy consumption is one of the major concerns of the research and design community, from circuit designers to embedded software developers. However, integrated system-level approaches which allow for an energy efficient mapping of an application on a customized platform are still lacking. What the research community is currently looking after are accurate system-level energy models, power estimation and analysis tools, functionality mapping and scheduling techniques, energy efficient communication synthesis and memory hierarchy optimization, as well as energy aware software compilation techniques. Hardware components that must be efficiently utilized include processors and memories. Highly optimized and low-cost processors can be designed with tools that support the creation of tuned micro-architectures and tool chains for application-specific instruction set processors. Moreover, highly efficient use of memories is increasing critical, both because the speed gap between processors and memories widens, and because the power consumed in memory systems is rapidly increasing.

One additional challenge which is taking center spot is the management of leakage power. Current nanometer technologies are said to be “leakage dominated”. Even though the amount of leakage power is still 3 orders of magnitute less than dynamic power, its incidence on battery lifetime is rapidly increasing for two main reasons. First, the complexity (and the number of transistors) of integrated circuits is growing, and an increasing number of idle functional unit and memories are present, at any given time on a given die, thereby contributing to an increased quiescent power. Moreover, leakage power is a strong function of temperature, which is an “exogenous” variable. Hence, the power consumed by a digital system is becoming a strong function of its operating conditions. This impacts in a significant way the techniques that should be used for power management, as decisions and policies should now depend on temperature.

Real-time applications, hard or soft, are raising the challenge of predictability. This is an extremely difficult problem in the context of modern, dynamic, multiprocessor platforms which, while providing potentially high performance, make the task of timing predictability extremely difficult (if at all solvable without drastically limiting the spectrum of applicable architectures). . 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. Software generation and performance evaluation tools have to be made aware of the particularities of a certain memory hierarchy, or the dynamic features of the processor micro-architecture, such as to be able to both generate efficient code and accurately predict performance numbers. The basic dilemma which researchers still face is how much to compromise predictability in order to improve average performance? Or how much can cost and average performance be affected in order to achieve a predictable system with good worst case behaviour? This always increasing inter-dependence between 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.

2.1.3 Simulation and Performance Analysis

The success of such new design methods depends on the availability of analysis techniques, beyond those corresponding to the state-of-the-art. Today, manufacturers and suppliers still rely only on extensive simulation to determine if the timing constraints are satisfied. However, simulations are very time consuming and provide no guarantees that imposed requirements are met.

There is a large quantity of research related to scheduling and schedulability analysis, with results having been incorporated in analysis tools available on the market. However, the existing approaches and tools address the schedulability of processes mapped on a single processor, or the schedulability of messages exchanged over a given communication protocol.

Several research groups have provided analyses that bound the communication delay for a given communication protocol, and extended the uni-processor analysis to handle distributed systems. However, none of the existing approaches offers an analysis that can handle applications distributed across different types of networks (e.g., CAN, FlexRay, TTP) consisting of nodes that may use different scheduling policies (e.g., static cyclic scheduling, fixed-priority preemptive scheduling, earliest deadline first).

Current modeling and design approaches often dimension systems for the worst case. However, embedded systems are growing more complex and dynamic than they used to be. E.g. in multi-media embedded systems, bit rates and encoding effort may vary by orders of magnitude depending on the complexity of the audio or video being played out, the complexity of the compression and on the required quality. Additionally, the embedded devices and application's functionality increases and they become more open to interaction with their environments. E.g. users may issue requests to the applications to change the resolution or frame rate. High-quality multimedia delivery on affordable embedded system's hardware requires cost-efficient realization of high throughput processing that is guaranteed to deliver the required performance. The platform needs to have sufficient resources to process the stream, even under the highest load conditions. Yet, it should not waste too many available resources when complexity of the stream is less. Current design approaches for multimedia embedded systems cannot deal appropriately with the increasing dynamism inside applications and the dynamically changing set of running applications. Often design approaches are based on worst-case analysis, resulting in over-dimensioned systems.

In order to (automatically) take informed design decisions, accurate analysis techniques are needed to:

- handle distributed applications, data and control dependencies;

- accurately take into account the details of the communication protocols;
- handle heterogeneous scheduling policies;
- take into account the fault-tolerance techniques used for dependability;
- capture the integration of control models and streaming models to understand the effects of sporadic events interfering with ongoing dataflow computations;
- efficiently use platform resources through statistical multiplexing based on a combination of worst-case and stochastic techniques.

2.2 Main Aims for Integration and Building Excellence through Artist2

Following the activities presented in the previous section, the cluster on execution platforms follows the following strategy and uses the following mechanisms to spread the knowledge and integration achieved so far:

- Summer Schools and Training Activities to distribute the knowledge acquired in ARTIST2 to (a) other countries, (b) other communities and (c) young researchers.
- Tutorials at major conferences to reach new and larger research communities.
- Joint publications between partners which not only show the integration within the cluster but are an excellent instrument to disseminate the integration results.
- New research projects with industrial partners which allow us to apply the obtained results at an industrial scale. This way, we also receive feedback and ideas for new research directions.
- Cooperation with other research groups, especially outside the EU (mostly USA and Asia). In this case, spredading exellenge is not the only objective. The cluster participants can be exposed to new research problems and new approaches, that can be then explored and improved within the cluster.

2.3 Other Research Teams

It appears that main research groups in Europe dealing with execution platforms for embedded systems are in the ARTIST2 network, either as full or as affiliated partners. There are some exceptions though, caused by the fact that not all are accepting a European network of Excellence as a viable funding instrument. In the following, some of these groups are listed together with their relation to ARTIST2.

IMEC's MPSoC research team. The MPSoC research team focus on platforms containing multiple heterogeneous, flexible processing elements, a memory hierarchy and I/O components. All these components are linked to each other by a flexible on-chip interconnect structure. The main focus of the team is to develop architectures which can meet the performance needs of multimedia applications, while limiting the power consumption. They are dealing with topics like design-time application exploration and optimization, platform architecture and runtime management. Close interaction, both at design time and at runtime, between these three topics creates a global solution that meets the MPSoC environment needs. In particular the work on the ADRES core, a platform containing a coarse-grained reconfigurable array, has been explored in a cooperation with researchers in ARTIST2, in which Jan Madsen, DTU, has participated.

The University of North Carolina at Chapel Hill, Sanjoy Baruah and Jim Anderson. Sanjoy Baruah and Jim Anderson are known in particular for their research in the domain of multiprocessor real-time scheduling.

University of Dresden, Hermann Härtig. Hermann Härtig is a leading researcher in the domain of micro-kernel based real-time operating systems.

Low power embedded systems design: In the area of low power embedded systems design, several new and relevant research thems are explored by other teams, not included in the ARTIST2 network. In particular research groups in the USA have a long tradition of excellence in low power research. We can mention the group lead by prof. Jan Rabaey in UC Berlely, which is carrying out ground-breaking work on hardware platforms for wireless sensor networks. In the same area, several other groups are performing top-level research, e.g. Anantha Chandrakasan's group at MIT and David Blaauw's group at University of Michigan. Low power execution platforms are not relevant only for wireless sensor network, but also for mobile computing and even for servers and traditional computing infrastructure (e.g. servers). In these areas, the groups lead by Profs. Vijaykrishnan Narayanan, Mahmut Kandemir and Mary Jane Irwin at Penn State University, has produced a large number of interesting results in the last few years. We mention in particular their work on power issues for 3D integration and their analysis of power vs. reliability tradeoffs in high-performance computing. In this area, very interesting work is also performed by the group of prof. Kevin Skadron. The focus of this group is on thermal issues, which are very significant for high-performance system.

Universita degli Studi di Verona/Electronic Design Automation (EDA) group, Prof. Franco Fummi. Main research activities of the EDA group concern system verification, system synthesis and optimization, hardware description languages, power consumption, language abstraction, and system testing. Interactions with members of the execution platforms cluster are, for example, by participation in European projects (e.g. the STRP "Vertigo") together with the Linköping group.

University of Southampton,/Electronic Systems Design Group, Prof. Bashir Al Hashimi. The Electronic Systems Design (ESD) Research Group is internationally recognized in two main areas - the development of novel algorithms and methodologies for Electronic Design Automation to support the design and test of large systems, and for intelligent sensor micro-systems. The group is working in the areas of system modeling, simulation, and synthesis, SoC design and testing, as well as smart sensors. Several cooperation projects have been undertaken, in particular with the Linköping group.

Carnegie Mellon University/System Level Design Group/Prof. Radu Marculescu. The System Level Design group performs research on formal methods for system-level design of embedded applications. They, in particular, focus on fast methods for power and performance analysis that can guide the design process of portable information systems. Important results have been obtained with regard to the communication-centric SOC design, providing formal support for analysis and optimization of novel on-chip communication architectures. In particular, this work addresses fundamental research problems for defining scalable and flexible communication schemes via the Network-on-Chip (NoC) approach. Interaction has been by, for example, PhD student exchanges with the Linköping group.

University of Southampton, Electronic Systems Design Group Prof. Bashir M. Al-Hashimi. The System design group is active in a number of activities related to execution platforms. The main activities of this group are the following: (i) Design methods and tools for large-scale integrated systems, with special emphasis on the issues raised by deep submicron technology; (ii) Next-genration interconnection technologies. Emphasis is placed upon the employment of the emerging concept of Network-on-Chip (NoC) proposed to overcome complex on-chip communication problems, where SoC cores communicate with each other using packets through interconnection network, thus providing support for communication infrastructure reuse, reliable and power efficient interconnection technology. (iii) Low-power built-in self-test. The focus is to investigate in detail some of the promising low power DFT techniques that have been recently developed. The availability of low power BIST techniques and architectures would allow IC designers to address concurrently design and test with the aim of generating self-testable designs that are not only optimised in terms of silicon area but also dissipate less power during test than in functional mode, hence resulting in safer testing.

2.4 Interaction of the Cluster with Other Communities

Two members of the cluster on Execution Platforms haven been given part of a summer school/advanced course on ADVANCED DIGITAL SYSTEMS DESIGN on 25.-29th September, Lausanne, Switzerland. The participants are from industry and university. This way, results from the integrated view of embedded system design will be brought to a much larger community.

A Workshop Models of Computation and Communication has been taken place atETH Zurich November 16th and 17th 2006. It brought together scientists from various areas, i.e. formal methods, hardware design and software architecture, see http://www.artist-embedded.org/artist/MoCC-06.html.

ETH Zurich has been organizing and participating in the CASTENESS Workshop, see www.casteness.org. The workhop put together the expertise of various EU projects such as ARTIST2, SHAPES, AETHER. In addition, ETH Zurich has been given a tutorial on issues that have been investigated in the ARTIST2 context: Analytic Performance Estimation, Mapping Algorithms to Architectures, Scalable SW Construction. The workshop has been sponsored by ARTIST2 and took place 15.-17th of January 2007.

ETH Zurich has been organizing a Workshop at a major conference in the area of Embedded Software (EMSOFT): "Foundations and Applications of Component-based Design", October 26th 2006, Seoul. The workshop has been organized in the framework of the Embedded Systems Week (http://www.esweek.org/), which federates CODES/ISSS, EmSoft, and CASES.

ETH Zurich has been organizing a Dagstuhl Seminar 04.03.2007-09.03.2007: "Quantitative Aspects of Embedded Systems". The purpose was to connect the results on performance analysis in ARTIST2 with the community dealing with statistical and stochastic methods. Therefore, organizers of this workshop have been B. Haverkort (Univ. of Twente, NL), J.-P. Katoen (RWTH Aachen, DE), L. Thiele (ETH Zürich, CH) see http://kathrin.dagstuhl.de/07101/.

ETH Zurich has been the general chair of the ARTIST2-sponsored conference ARCS'07: "Architecture of Computing Systems", that took place at the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland, March 12-15, 2007, http://arcs07.ethz.ch/. Here, a broad audience was present which allowed disseminating results on embedded system design methods to a larger community.

TU Eindhoven has organised and given a tutorial on system-level performance modelling at the fifth ACM-IEEE International Conference on Formal Methods and Models for Codesign (MEMOCODE'2007), May 30 - June 1st, Nice, France.

TU Eindhoven has organised a workshop on Models of Computation and their application in the design of multi-processor systems. The workshop tool place in Eindhoven, May 16, 2007, see http://www.es.ele.tue.nl/premadona/moc.php.

TU Eindhoven (Twan Basten) has co-organized ACSD 2007, the 7th International Conference on Application of Concurrency to System Design. ACSD was held in Bratislava, Slovak Republic, 10-13 July 2007. Lothar Thiele gave an invited presentation at the conference, on Modular Performance Analysis applied to System-on-Chip application mapping.

DTU has been the program chair for the DATE'07 conference, "Design, Automation, and Test in Europe", that took place in Nice, France, April 16-20, 2007. http://www.date-conference.com/conference/2007/prog/ . Many ARTIST2 partners were involved in organizing tracks and topics related to embedded systems design.

DTU has organized an ARTIST2 sponsored PhD course on "Advanced Topics in Embedded Systems", that took place at IMM, DTU, Lyngby, Denmark, June 4-12, 2007. Lectures were given by ARTIST memebers from Oldenburg, Germany, ETH Zurich, Switzerland, and Braunsweigh, Germany. http://www.artist-embedded.org/artist/ARTIST2-PhD-Course-on-Automated.html

DTU has been co-organizing a workshop on "Tool Platforms for Embedded System Modelling, Analysis and Validation", which took place in Berlin, Germany, July 1-2, 2007. The workshop was arranged as a satelite event to CAV 2007, "Computer Aided Verification". http://www.artist-embedded.org/artist/Aims-and-Scope.html

DTU has been given part of a PhD seminar (Trust Soft), which tok place in Oldenburg, Germany, one week in July 2007.

Linköping has given a tutorial at the "International Workshop on Embedded Systems 2006", Seul October 2006. With this occasion several results obtained in the ARTIST context have been made accessible to an international audience.

TU Braunschweig has organized together with ETH Zürich and University of Notre Dame the tutorial "Extensible Frameworks for System-Level Analysis of Real-Time Systems" at the Real-Time and Embedded Technology and Applications Symposium (RTAS). The tutorial took place April 4, 2006.

TU Braunschweig participated in the "Workshop on Models and Analysis for Automotive Systems" at the Real-Time Systems Symposium (RTSS). The talk was named and discusses "The Need of a Timing Model for the AUTOSAR Software Standard". The workshop took place December 5, 2006.

TU Braunschweig has been organizing the Embedded Software Track at the major European conference on design automation DATE (Design Automation and Test in Europe) that took place April 16-20, 2006. The track was devoted to modelling, analysis, design and deployment of embedded software, including formal methods, tools, methodologies and development environments. Thereby, the emphasis was on embedded software platforms, software integration and portability issues.

TU Braunschweig has given a lecture with the title "Supporting Predictable Design Using Formal Analysis Techniques" at the ARTES summerschool (A Network for Real-Time Research and Graduate Education in Sweden) that took place in Nässlingen, Sweden, August 23, 2006. The audience consisted of Phd Students from the field of real-time research, which allowed disseminating recent results in embedded system design to related research teams (mainly) in Scandinavia.

TU Braunschweig was invited to participate in the special session on "Virtual Automotive Platforms" at the renowned Design Automation Conference (DAC). The talk "Automotive Software Integration" showed how formal techniques can be applied to solve performance related integration problems in the design process of complex modern automotive systems. The special session took place June 6, 2007.

TU Braunschweig was invited to participate in the ARTIST workshop on "Tool Platforms for Modelling, Analysis and Validation of Embedded Systems" at the conference on Computer Aided Verification (CAV).The talk "SymTA/S - Modeling system timing using abstract event streams" allowed disseminating results in the field of compositional performance verification techniques to a larger community. The workshop took place July 1-2, 2007.

3. Overall Assessment and Vision for the Cluster

3.1 Assessment

There has been substantial progress in integrating different research directions and view points. Indicators that show this clearly are (a) the joint participation in summer schools, workshops and tutorials and (b) the number and quality of joint publications, and (c) the integration of tools.

3.1.1 Integration Activities within the Cluster

The following list summarizes the integration activities during the last year within the cluster on execution platforms:

- *ETHZ-TUBS:* Combination of the MPA and Symta/S methods by providing a tool interconnection. This result was achieved by mutual visits and implementation effort. The results are published in [25].
- ETHZ-DTU: Participation of ETH Zurich in the PhD course at DTU on formal methods for embedded systems. June 4 – 12 2007, Nikolay Stoimenov, Simon Perathoner, Wolfgang Haid. Title: System Level Performance Anlysis Using MPA-RTC: Models, Methods and Scenarios, see http://www.artist-embedded.org/artist/ARTIST2-PhD-Course-on-Automated.html.
- ETHZ-TUBS-Universidad Cantabria: Comparison of various methods for the performance analysis of distributed embedded systems. The results are based on a set of benchmarks that have been constructed during an ARTIST2 workshop in Leiden, organized by ETHZ. The results of the comparison are published in [26].
- ETHZ-BOLOGNA: There have been major results by the combination of expertise in Bologna and Zurich in the area of energy harvesting systems. Davide Brunelli from Bologna visited ETH Zurich several months in order to combine the existing expertise in hardware, software, low power and real-time scheduling. As a result of this acitivity, two publications emerged: [27], [28].
- LINKOEPING-BOLOGNA: a detailed assessment of allocation and scheduling techniques for low-power MPSoC has been performed. More specifically, a heuristic allocation technique based on genetic algorithms developed in LINKOEPING has been compared with an exact optimization approach developed in Bologna. This has been achieved thanks to the joint work performed by A. Andrei in LINKOEPING, and C. Nardini, M. Ruggiero, A. Guerri in UNIBO. As a result of this activity one joint publication emerged, [29], and one additional one has been submitted for publication.
- TU/e-Leuven. There is an ongoing cooperation between the TU/e and IMEC/KU Leuven, to develop scenario-based design. Traditional use-case scenarios focus on the user-system interaction, whereas this focuses on system scenarios that classify system behaviors from the resource perspective. The use of system scenarios allows for more aggressive design decisions per system scenario. See http://www.es.ele.tue.nl/scenarios.
- ETHZ-TU/e. Valentin Gheorghita, a senior PhD student of the TU/e, visited the group of Lothar Thiele for a few weeks during the summer of 2006, to explore possibilities for cooperation.

- DTU-LiU: Proposed an approach to the scheduling and voltage scaling of embedded real-time applications that decides the voltage levels and start times of processes and the transmission times of messages, such that the transient faults are tolerated, the timing constraints of the application are satisfied and the energy consumption in the no-fault scenario is minimized. The results are published in [30], [31].
- DTU-LiU: Timing analysis of applications communicating over a FlexRay bus and optimization approach for determining a FlexRay bus configuration which is adapted to the particular features of an application and guarantees that all time constraints are satisfied. The results are published in [32], [33].
- DTU-IMEC: Extending the ADRES architecture and toolchain to support multitask programming. ADRES is an architecture for coarse grained dynamic reconfiguration. The results of the cooperation are published in [35].
- DTU-BOLOGNA: The joint work on traffic generators for network-on-chip simulation has been completed with 3 additional publication [36], [37], and [38].
- Linköping- DTU: The groups at Linköping and DTU have closely cooperated on the topic of fault tolerant real-time systems. Viaceslav Izosimov, PhD student at Linköping has visited DTU for several weeks. Results are published in [29], [30], [31].
- Linköping-Bologna: Linköping and Bologna have a cooperation on two topics: (1) Energy optimisation for multiprocessor systems: Task scheduling, mapping and voltage selection techniques using mathematical programming techniques has been investigated. Results are published in [24]. (2) Predictability on multiprocessor systems: Paolo Burgio, a master student from Bologna, working on this topic, is visiting Linköping for several months.
- Linköping-Braunschweig: The approach to predictable implementation of real-time systems on multiprocessor platforms, developed at Linköping, is using for WCET analysis a technique that has been implemented on top of the Symta/P tool from Braunschweig.
- TUBS-DTU: TU Braunschweig participated in the PhD course on “Formal Methods for Embedded Systems” at the Technical University of Denmark (DTU). Arne Hamann and Razvan Racu held a full day tutorial with the title “Symbolic Timing Analysis for Systems (SymTA/S)” covering the formal foundations of the SymTA/S approach as well as advanced analysis features such as sensitivity analysis, design space exploration, and robustness optimization techniques. During hands-on sessions the presented formal techniques were applied to solve practical performance verification problems. The tutorial took place June 12, 2007. Web link: http://www.artist-embedded.org/artist/ARTIST2-PhD-Course-on-Automated.html.

3.1.2 List of Joint Publications

The following list contains publications, where authors are in different research sites which are participating in the ARTIST2 network and where at least one author is in the cluster on Execution Platforms. It clearly shows the degree of integration that has been achieved. The following list collects all joint publications since the start of ARITST2:

[1] Lothar Thiele, Reinhard Wilhelm: Design for Timing Predictability. Dagstuhl Online, 2004.

[2] Lothar Thiele, Reinhard Wilhelm: Design for Timing Predictability. Journal on Real Time Systems, 2004.

[3] S. Mahadevan, F. Angiolini, M. Storgaard, R. G. Olsen, J. Sparsø, and J. Madsen. A network traffic generator model for fast networkon-chip simulation. In Proceedings of Design, Automation and Testing in Europe Conference 2005 (DATE05)

[4] F. Angiolini, S. Mahadevan, J. Madsen, L. Benini, and J. Sparsø. Realistically rendering soc traffic patterns with interrupt awareness. In IFIP International Conference on Very Large Scale Integration (VLSI-SoC), September 2005.

[5] Ernesto Wandeler, Lothar Thiele, Marcel Verhoef, Paul Lieverse: System Architecture Evaluation Using Modular Performance Analysis - A Case Study. Paphos, Cyprus, October, 2004.

[6] Paul Pop, Petru Eles, Zebo Peng, Viacheslav Izosimov, Magnus Hellring, Olof Bridal: Design Optimization of Multi-Cluster Embedded Systems for Real-Time Applications. I Design, Automation and Test in Europe (DATE 2004).

[7] M. Loghi, F. Angiolini, D. Bertozzi, L. Benini, R. Zafalon, "Analyzing on-chip communication in a MPSoC environment", IEEE/ACM Design, Automation and Test in Europe, Paris, France, Feb 2004.

[8] Jan Staschulat, Rolf Ernst, Andreas Schulze, Fabian Wolf. Context Sensitive Performance Analysis of Automotive Applications. In Designer's Forum at Design, Automation and Test in Europe (DATE), 2005.

[9] Bren C. Mochocki, Xiaobo S. Hu, Razvan Racu, Rolf Ernst. Dynamic Voltage Scaling for the Schedulability of Jitter-Constrained Real-Time Embedded Systems, ICCAD 2005.

[10] E. Wandeler, L. Thiele, M. H. G. Verhoef, P. Lieverse. System Architecture Evaluation Using Modular Performance Analysis - A Case Study. Journal Software Tools for Technology Transfer (STTT), 2006.

[11] Simon Künzli, Francesco Poletti, Luca Benini, Lothar Thiele: Combining Simulation and Formal Methods for System-Level Performance Analysis, IEEE Design Automation & Test in Europe (DATE), Munich, Germany, March 2006.

[12] C. Moser, D. Brunelli, L. Thiele, and L. Benini, “Real-time scheduling with regenerative energy.”, 18th Euromicro Conference on Real-Time Systems (ECRTS 2006), Dresden, Germany, July 5-7, 2006.

[13] C. Moser, D. Brunelli, L. Thiele, and L. Benini, “Lazy scheduling for energy harvesting sensor nodes.”, The Fifth IFIP Working Conference on Distributed and Parallel Embedded Systems (DIPES 2006), Braga, Portugal, October 13-15, 2006.

[14] Bren Mochocki, Xiaobo Sharon Hu, Razvan Racu, Rolf Ernst. Dynamic Voltage Scaling for the Schedulability of Jitter-Constrained Real-Time Embedded Systems. In *International Conference on Computer Aided Design (ICCAD)*, San Jose, USA, November 2005.

[15] Razvan Racu, Arne Hamann, Rolf Ernst, Bren Mochocki, Sharon Hu, Methods for Power Optimization in Distributed Embedded Systems with Real-Time Requirements, In Proc. International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES), Seoul, Korea, October 2006

[16] Kai Richter, Marek Jersak, Rolf Ernst. How OEMs and suppliers can tackle the network dimensioning problem. *Embedded Real Time Software Congress (ERTS06)*, Toulouse, France, January 25-27, 2006.

[17] Kai Richter, Rolf Ernst. Real-Time Analysis as a Quality Feature: Automotive Use-Cases and Applications, *Embedded World 2006 Fair and Conference*. Nuremberg, Germany - February 14-16, 2006.

[18] Kai Richter, Marek Jersak, Rolf Ernst. How OEMs and suppliers can face the network dimensioning challenges. *Design, Automation and Test in Europe (DATE) Conference,* Special Track Automotive Designer's Forum, Munich, Germany, March, 2006.

[19] Kai Richter, Rolf Ernst. Applying Real-Time Network Research in the Automotive Industry: Lessons Learned and Perspectives, *Euromicro Conference on Real-Time Systems (ECRTS), s*atellite workshop on Real Time Networks (RTN)*,* Dresden, Germany, July 2006.

[20] Kai Richter, Marek Jersak, Arne Hamann, Rolf Ernst. Scheduling Analysis in the Automotive Design Flow. *ARTIST2 Workshop at the International Conference on Embedded Sofware (EMSOFT),* Seoul, Korea, October, 2006.

[21] L. A. Cortes, P. Eles, Z. Peng, "Quasi-Static Scheduling for Multiprocessor Real-Time Systems with Hard and Soft Tasks", 11th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA'05), Hong Kong, 2005.

[22] L. A. Cortes, P. Eles, Z. Peng, "A Quasi-Static Approach to Minimizing Energy Consumption in Real-Time Systems under Reward Constraints", 12th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA'06), Sydney, Australia, 2006.

[23] L. A. Cortes, P. Eles, Z. Peng, "Quasi-Static Assignment of Voltages and Optional Cycles in Imprecise-Computation Systems with Energy Considerations", IEEE Trans. on Very Large Scale Integration Systems, Vol. 14, Issue 10, October 2006, pp. 1117-1129

[24] M. Ruggiero, G. Pari, A. Guerri, M.Milano, D. Bertozzi, L.Benini, A Andrei, "A Cooperative, accurate solving framework for optimal allocation, scheduling and frequency selection on energy-efficient MPSoCs", International Symposium on System-on-Chip, Tampere, Finland, 2006.

[25] Simon Künzli, Arne Hamann, Rolf Ernst, Lothar Thiele: Combined Approach to System Level Performance Analysis of Embedded Systems CODES/ISSS 2007, Salzburg, Austria, 2007.

[26] S. Perathoner, E. Wandeler, L. Thiele, A. Hamann, S. Schliecker, R. Henia, R. Racu, R. Ernst, M. González Harbour. Influence of Different System Abstractions on the Performance Analysis of Distributed Real-Time Systems. In 7th ACM Conference on Embedded Systems Software (EMSOFT), Salzburg, Austria, October, 2007.

[27] C. Moser, L. Thiele, D. Brunelli, and L. Benini, "Adaptive Power Management in Energy Harvesting Systems.", in Design, Automation and Test in Europe (DATE),Nice, France, April 2007.

[28] C. Moser, D. Brunelli, L. Thiele, and L. Benini, "Real-time scheduling for energy harvesting sensor nodes.", to be published in Real-time Systems, Special Issue on Real-Time Wireless Sensor Networks.

[29] V. Izosimov, P. Pop, P. Eles, Z. Peng, "Synthesis of Fault-Tolerant Schedules with Transparency/Performance Trade-offs for Distributed Embedded Systems", in Proceedings of DATE, 706–711, 2006.

[30] P. Pop, K. H. Poulsen, V. Izosimov, P. Eles, "Scheduling and Voltage Scaling for EnergyReliability Trade-offs in Fault-Tolerant Time-Triggered Embedded Systems", ACM/IEEE International Conference on Hardware-Software Codesign and System Synthesis, 2007.

[31] K. H. Poulsen, P. Pop, V. Izosimov, “Energy-Aware Synthesis of Fault-Tolerant Schedules for Real-Time Distributed Embedded Systems”, Euromicro Conference on Real-Time Systems (ECRTS, WiP session), 2007

[32] T. Pop, P. Pop, P. Eles, Z. Peng, A. Andrei, “Timing Analysis of the FlexRay Communication Protocol”, Proceedings of the 18th Euromicro Conference on Real-Time Systems (ECRTS 06), 203-213, 2006.

[33] T. Pop, P. Pop, P. Eles, Z. Peng, “Bus Access Optimisation for FlexRay-based Distributed Embedded Systems”, Proceedings of the Design, Automation, and Test in Europe Conference (DATE'07), 51-62, 2007.

[34] M. Fränzle, M. R. Hansen, *Deciding an Interval Logic with Accumulated Durations*, TACAS 2007, LNCS 4424, pp. 201-215, Springer-Verlag, 2007

[35] Kehuai Wu, Andreas Kanstein, Jan Madsen and Mladen Berekovic *MT-ADRES: Multithreading on Coarse-Grained Reconfigurable Architecture* International Workshop on Applied Reconfigurable Computing 2007.

[36] Mahadevan, S., Angiolini, F., Sparsø, j., Benini, L., Madsen, J., *A Cycle-True Traffic Generator Capturing Context Switching Effects*, IFIP VLSI SoC: From Systems to Chip, IFIP, 2007. Accepted for publication.

[37] Mahadevan, S., Angiolini, F., Sparsø, J., Benini, L., Madsen, J., A Reactive IP Emulator for Multi-Processor System-on-Chip Exploration, *IEEE Transaction on CAD*, 2007. Accepted for publication.

[38] Mahadevan, S., Angiolini, F., Storgaard, M., Olsen, R., Sparsø, J., Madsen, J., A Network Traffic Generator Model for Fast Network-on-Chip Simulation, In “Design, Automation, and Test in Europe: The Most Influencial Papers of 10 Years DATE”, R. Lauwerein, J. Madsen, editors. Sppringer, 2007. Accepted for publication.

[39] Ruggiero M., Pari G., A. Guerri, D. Bertozzi, M. Milano, L. Benini, Andrei A., “A Cooperative, Accurate Solving Framework for Optimal Allocation, Scheduling and Frequency Selection on Energy-Efficient MPSoCS”, Proceedings of The IEEE International SOC Conference (SOCC), 2006.

[40] C. Couvreur, V. Nollet, F. Catthoor and H. Corporaal, Fast Multi-Dimension Multi-Choice Knapsack Heuristic for MP-SoC Run-Time Management. In Proceedings of the International Symposium on System-on-Chip (SOC), pp. 195-198, 2006.

[41] T. Vander Aa, M. Jayapala, F. Catthoor, G. Deconinck and H. Corporaal, Instruction Transfer And Storage Exploration for Low Energy VLIWs. In proceedings of IEEE 2006 Workshop on Signal Processing Systems (SiPS 2006), 2006, pp. 313-318.

[42] K. Denolf, A. Chirila-Rus, P. Schumacher, R. Turney, K. Vissers, D. Verkest, and H. Corporaal, A Systematic Approach to Design Low Power Video Codec Cores. Accepted by EURASIP Journal on Embedded Systems.

[43] V. Nollet, P. Avasare, H. Eeckhaut, D. Verkest and H. Corporaal, Run-Time Management of a MPSoC Containing FPGA Fabric Tiles, Accepted by Special Issue TVLSI.

[44] C. Couvreur, V. Nollet, T. Marescaux, E. Brockmeyer, F. Catthoor and H. Corporaal, Design-time application mapping and platform exploration for MP-SoC customized run-time management, To be published in IEE Proceedings Computers & Digital Techniques, 2007.

[45] T. Marescaux and H. Corporaal, Introducing the SuperGT Network-on-Chip. In proceedings of DAC, 2007.

[46] P. Raghavan, A. Lambrechts M. Jayapala, F. Catthoor, D. Verkest and H. Corporaal, Very wide register: an asymmetric register file organization for low power embedded processors. In proceedings of DATE 2007.

[47] T. Pop, P. Pop, P. Eles, Z. Peng, “Bus Access Optimisation for FlexRay-based Distributed Embedded Systems”, Design, Automation, and Test in Europe Conference (DATE'07), Nice, France, April 2007, pp. 51-62.

[48] V. Izosimov, P. Pop, P. Eles, Z. Peng, “Mapping of Fault-Tolerant Applications with Transparency on Distributed Embedded Systems”, 9th Euromicro Conference on Digital System Design (DSD 2006), Dubrovnik, Croatia, September 2006, pp. 313-320.

[49] P. Pop, K. H. Poulsen, V. Izosimov, P. Eles, “Scheduling and Voltage Scaling for Energy/Reliability Trade-offs in Fault-Tolerant Time-Triggered Embedded Systems”, CODES/ISSS 2007, Salzburg, Austria, 2007.

3.2 Vision and Long Term Goals

The research in embedded systems still is fragemented. This not only is true within a single subject but also between several subdisciplines. Examples are the parallel efforts in real-time scheduling and real-time analysis in the area of ‘Execution Platforms’, ‘Hard Real-Time Systems’ and ‘Software Components’. It is one of the major goals of the cluster on ‘Execution Platforms’ to establish closer links to the other communities and to take advantage of the scientific results and insights.

Cross-layer design is a key issue in embedded systems. The classical view of a strict layering according to chosen abstraction levels does not work any more because of the importance of non-functional constraints and limited resources. Therefore, completely new concepts are necessary that enable the integrated modelling and design under predictability AND efficiency constraints. It is expected that this move towards a resource-aware design trajectory involves all current layers and a breakthrough can be obtained by integration only.

3.3 Plans for Year 4

<u>System Modelling Infrastructure</u>

Within the system modelling infrastructure, we are planning the following activities within the next 12 months:

- DTU and LiU will continue its work on models for the analysis and optimization of fault-tolerant embedded systems. In particular, they will refine the model to capture fine-grained combinations of several fault-tolerance techniques. At the moment, replication and re-executions can not be combined in the FT-PG model for a single process. There are several ways such a situation can be captured, and they will evaluate the advantages and disadvantages of each approach.
- TU Braunschweig will continue its work extending the semantic model of SymTA/S to efficiently cover MPSoC architectures. Additionally, they will further investigate challenging timing issues in multiprocessor systems. Main goals include increasing the precision of latency prediction for multiple memory accesses, and looking into the issue of resource synchronization.
- DTU will continue the work on formalizing the ARTS model using timed automata based on UPPAAL. This will allow the same platform model to be expresed as a simulation model and as a formal model. The work will be carried out in cooperation with the

research group at CISS in Aalborg which is a partner in the ARTIST2 cluster on test and verification. Work on extending the simulation model of ARTS will be continued.

- DTU will refine its formal model to address modeling and verification issues closer to the hardware layer of the execution platform.
- ETHZ intends to combine Modular Performance Analysis with timed automata based evaluation methods. This work will be done togehter with the affiliated parter NUS (National University Singapore). This way, there is a link and integration of MPA with (a) simulation (done in a joint work togehter with University Bologna), (b) Symta/S (joint work with TU Braunschweig) and (c) timed automata.
- LiU will do experimental evaluations using the simulation environment for distributed embedded systems: of particular interest are issues related to fault tolerance.

Communication-Centric Systems

ETH Zurich plans to build on the results of the previous year in terms of comparing different analysis methods in terms of scope and accuracy. As a result, we expect to combine even more different formalisms, e.g. timed automata, and to improve the properties of our own method, Modular Performance Analysis.

The Linköping group will continue further development of the analysis and optimization techniques for fault-tolerant and predictable distributed systems. In particular, fault tolerance for soft real-time systems will be investigated. The approach to predictable implementation of real-time systems on multiprocessor SoC will be further developed.

DTU and LiU will continue their collaboration on the FlexRay communication protocol. So far, the topic of transmission reliability has not been addressed during FlexRay optimization. FlexRay has two independent channels that can be used for message replication to increase redundancy. Another option is to retransmit the message in order to protect against faults. DTU and LiU will determine reliability equations for FlexRay messages, and propose an optimization method to increase reliability of transmission. The optimization will decide if a message is to be replicated on the two channels, or retransmitted, such that the timing constraints are still satisfied.

DTU will continue its work on efficient NoC architectures. In particular the will work on techniques to automatically generate efficient NoC topologies. They will investigate reconfigurable NoC structures based on combining circuit switching and packet switching. Finally, work will be started on linking the NoC architecture with the application program through the development of programming models for NoC-based systems.

TU Braunschweig will further investigate the application of hierarchical event models for performance verification of embedded systems. One promising application will be the detailed analysis of effect that COM-Layers have on communication timing. Additionally, TU Braunschweig will pursue its research in the domain of online performance verification. Main focus will be the development of heuristic optimization algorithms performing what-if analyses to continously improve system robustness. Similar strategies will be investigated to deal with faulty hardware, such as degrading processors or applications breaking their service contract, in order to remain in a feasable system state.

Low-Power Design

- Scheduling based energy optimization for energy-scavenging wireless sensor
    - By implementing the algorithms developed at ETH Zurich on the solar scavenger prototype developed at the University of Bologna, it is planned to test the theoretical energy harvesting framework and demonstrate sustainable

operation using solar energy. To this end, an interface is needed which abstracts the physical characteristics of the energy conversion for the above software layer. First measurements of the physical energy conversion/storage process on the prototype revealed that a correction of the energy flow models used so far will be necessary.

  - By deploying sensor nodes – which are powered by solar energy – it is planned to demonstrate the usefulness of the theoretical results in a practical application, possibly in an outdoor setting.

- System-level resource allocation and scheduling – the following issues will be addressed:
  - Power optimization for nanometer platforms. Models should be extended to deal with limited number of operating points, leakage power, variations. Results from the activity “Power optimization and analysis for Nanometer technologies” will be exploited.
  - Mapping and scheduling for general task graphs, including complete dataflow graphs as well as models for non-deterministic behaviour (e.g. conditional task graphs)
  - Comparison between exact/complete optimizers and heuristic/incomplete optimizers will be undertaken to assess: the optimality gap of heuristic techniques and the run-time advantages obtainable by giving up optimality.
  - Temperature aware energy optimisation: A methodology for taking into consideration the temperature aspect during energy optimisation by voltage selection will be investigated.

<u>Resource-Aware Design</u>

Plans for year 4 in his activity are focused on the continuation of ongoing activities. More specifically, Bologna and Dortmund will continue cooperation on code generation for resource-awared embedded platforms, with special emphaisis on multi-core architectures. Problem related to the management of multiple scratchpad memories and on the parallelization of loops will be tackled.

Collaboration of DTU and Linköping will focus on adaptivity-related aspects, which will allow system reconfiguration in case of failures or changes in the environment. Linköping will address quasi-static fault-tolerant scheduling for mixed soft/hard time-triggered systems. DTU will address schedulability analysis for mixed hard/soft applications mapped on event-driven real-time systems. Comparison and integration of the two approaches is foreseen

At Linköping, work on "Predictability for Multiprocessor SoC Architectures" will be continued. Main goals include further optimization of the bus access and Controller design and synthesis. Bologna will also be involved in this work, focusing on optimizing bus controller implementation.

An increase in cooperation and joint activities is expected. Thanks to the pivotal role of ARTIST2, several new projects will be initiated, as they have successfully obtained FP7 funding: We mention here, projects: MIMEE and PREDATOR, which will start early in 2008.

4. Cluster Participants

4.1 Core Partners

Cluster Leader	
	Lothar Thiele (ETH Zurich)
Technical role(s) within Artist2	Main areas of research: Embedded Systems and Software Artist2 activities and role: Communication Centric Systems: Formal Performance Analysis, Linking Simulation and Verification, Design Space Exploration of Embedded Systems
Research interests	Research interests include models, methods and software tools for the design of embedded systems, embedded software and bioinspired optimization techniques.
Awards / Decorations	In 1986 he received the "Dissertation Award" of the Technical University of Munich, in 1987, the "Outstanding Young Author Award" of the IEEE Circuits and Systems Society, in 1988, the Browder J. Thompson Memorial Award of the IEEE, and in 2000-2001, the "IBM Faculty Partnership Award". In 2004, he joined the German Academy of Natural Scientists Leopoldina. In 2005-2006, he was the recipient of the Honorary Blaise Pascal Chair of University Leiden, The Netherlands.

Activity Leader for NoE Integration: Resource-aware Design	
	Luca Benini (University of Bologna)

Technical role(s) within Artist2	Main areas of research: Execution platforms Other projects involved in: low power design, communication-centric architectures Artist2 activities and role: JPRA Communication centric Systems, JPRA Design for Low Power (activity leader), JPRA Resource-aware design (co-activity leader)
Research interests	Research interests are in computer architecture and computer-aided design of digital systems, with special emphasis on low-power applications and SoC design.
Notable past projects	IST-Clean Contridbutes to the development of low-power design technologies for deep submicron technologies http://clean.offis.de/ MPARM project A collaborative infrastructure for MPSOC research. Virtual simulation platform in SystemC, with a complete software environment http://www-micrel.deis.unibo.it/sitonew/research/mparm.html
Awards / Decorations	IEEE Fellow
Further Information	Visiting professor at EPFL

Activity Leader for Cluster Integration: Communication-centric systems	
	Rolf Ernst (TU Braunschweig)
Technical role(s) within Artist2	Main areas of research: Embedded Systems Artist2 activities and role: Communication Centric Systems: Formal Performance Analysis, Design Space Exploration of Embedded Systems, Sensitivity Analysis of Embedded Systems. System modelling infrastructure: Flexibility Optimization of Embedded Systems.
Research interests	Research interests include embedded architectures, hardware-/software co-design, real-time systems, and embedded systems engineering.

Activity Leader for Platform: System Modelling Infrastructure	
	Jan Madsen (Technical University of Denmark)
Technical role(s) within Artist2	Main areas of research: Embedded Systems Design and Modeling Artist2 activities and role: System Modelling Infrastructure, Communication-Centric, Systems, Design for Low-Power
Research interests	Research interests include high-level synthesis, hardware/software codesign, System-on-Chip design methods, and system level modeling, integration and synthesis for embedded computer systems.
Role in leading conferences/journals/etc in the area	He is Program Chair for DATE07. He has been Tutorial Chair and Program Vice Chair for DATE06, Workshop Chair for CODES+ISSS'05, General Chair of CODES '01 and Program Chair of CODES '00. He is on the editorial board of the journal "IEE Proceedings – Computers and Digital Techniques"
Awards / Decorations	In 1995 he received the Jorck's Foundation Research Award for his research in hardware/software codesign

Activity Leader for Cluster Integration: Design for Low Power	
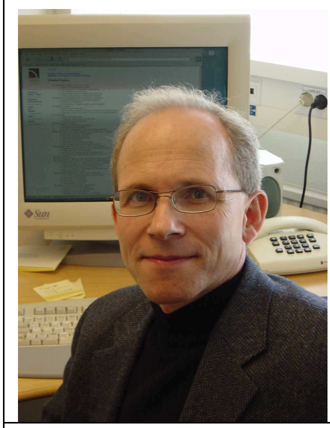	Petru Eles (Linköping University)
Technical role(s) within Artist2	Main areas of research: Embedded Systems Artist2 activities and role:Communication Centric Systems: analysis, optimisation. Low Power: modeling, optimisation, power

	management. System modeling infrastructure: modelling and simulation of distributed embedded systems.
Research interests	Research interests include electronic design automation, hardware/software co-design, real-time systems, design of embedded systems and design for testability.
Role in leading conferences/journals/etc in the area	- Associate Editor, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems; - Associate Editor, IEE Proceedings - Computers and Digital Techniques; - TPC Chair and General Chair, IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES/ISSS). - Topic chair, Design Automation and Test in Europe (DATE). - Topic Chair, Int. Conference on Computer Aided Design (ICCAD). - Program chair of the Hw/Sw Codesign track, IEEE Real-Time Systems Symposium (RTSS). - TPC Chair IEEE Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia). - Steering Committee Chair, IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES/ISSS).
Awards / Decorations	- Best paper award, European Design Automation Conference (EURO-DAC), 1992. - Best paper award, European Design Automation Conference (EURO-DAC), 1994. - Best paper award, Design Automation and Test in Europe (DATE), 2005. - Best presentation award, IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES/ISSS), 2003. - IEEE Circuits and Systems Society Distinguished Lecturer, for 2004 - 2005.

	Jeroen Voeten (assistant professor at the Information and Communication Systems group, faculty of Electrical Engineering of the Eindhoven University of Technology and research fellow at the Embedded Systems Institute in Eindhoven)
Technical role(s) within Artist2	Main areas of research: formal techniques for the specification, design and implementation of hardware/software systems. Artist2 activities and role: Performance Analysis in the System Design Process
Research interests	Research interests: formal specification and compositional verification of distributed real-time and fault-tolerant systems. Design of technical applications.

4.2 *Affiliated Industrial Partners*

	Magnus Hellring (Volvo Technology Corporation, Manager Systems and Architecture)
Technical role(s) within Artist2	Architecture and Design of Automotive Embedded Systems

	Roperto Zafalon (STM, Advanced System Technology – Research and Innovation, Manager Low Power System Design)
Technical role(s) within Artist2	Main areas of research: Low Power System Design Artist2 activities and role: JPRA design for low power, JPRA Resource-aware design

	Fabian Wolf (Volkswagen AG)
Technical role(s) within Artist2	Software and Hardware for Embedded Systems in Automotive Applications

	Kai Richter (SymTAVision GmbH)
Technical role(s) within Artist2	Formal Performance Analysis and Optimization of Embedded Systems, Reliable System Integration, Early Architecture Exploration

	Dirk Ziegenbein (Robert Bosch AG)
Technical role(s) within Artist2	Automotive Software Architectures

	Peter Mårtensson (Nokia Denmark)
Technical role(s) within Artist2	Platform architectures for mobile terminals

	Matthias Gries (Intel Germany)
Technical role(s) within Artist2	Microprocessor Technology Lab, new computer architecture for embedded systems

4.3 Affiliated Academic Partners

	Axel Jantsch (Royal Institute of Technology Stockholm)
Technical role(s) within Artist2	Design Methodology for Embedded Systems

	Giovanni DeMicheli (EPFL Lausanne)
Technical role(s) within Artist2	Main areas of research: Design Methodology for Embedded Systems, Low Power Design Artist2 activities and role: JPRA communication centric design

	Donatello Sciuto (Politecnico Di Milano)
Technical role(s) within Artist2	Design Methodology for Embedded Systems, Low Power Design

	Ed Deprettere (University Leiden, The Netherlands)
Technical role(s) within Artist2	Main areas of research: Design Methodology for Embedded Systems, Signal and Image Processing, Algorithm Design and Mapping Artist2 activities and role: Preparation of ARTIST2 Workhsop on Embedded Systems in Leiden, The Netherlands; Visists to Zurich in order to discuss on JPRA "Communication Centric Systems".

	Eugenio Villar, Pablo Sanchez (Universidad de Cantabria)
Technical role(s) within	Design and Implementation of Embedded H/S Systems

Artist2	

	Geert Deconinck (Katholiecke Universiteit Leuven)
Technical role(s) within Artist2	Dependability, embedded systems, control & automation, real-time, robust communication, interdependencies, critical information infrastructure protection

	Juergen Teich (University of Erlangen-Nuremberg)
Technical role(s) within Artist2	Expertise in FPGA technology and dependable system design.

	Luciano Lavagno (Politechnico di Turino)
Technical role(s) within Artist2	Asynchronous Circuit Design and Testing, H/S Codesign

4.4 Affiliated International Partners

	Sharon Hu (University of Notre Dame)
Technical role(s) within Artist2	Design for Low Power

5. Internal Reviewers for this Deliverable