IST-004527 ARTIST2
Network of Excellence
on Embedded Systems Design

Activity Progress Report for Year 3

JPRA-Cluster Integration

Design for Low Power

Clusters:

Execution Platforms

Activity Leader:

Petru Eles (Linköping University)
http://www.ida.liu.se/~petel

Policy Objective (abstract)

Power dissipation has become one of the most serious obstacles in the evolution of electronic systems. Recent developments in CMOS technology have resulted in increased power density, as well as increased active-state and stand-by power consumption. Furthermore, mainstream architectural design is moving towards energy-hungry programmable/configurable architecture meeting ever-increasing performance requirements. It is the objective of this activity to develop, promote and integrate methods that address power and energy issues across several layers of abstraction.

Table of Contents

1. Overview of the Activity

1.1 ARTIST Participants and Roles

Team Leader: Luca Benini – University of Bologna (Italy)

(i) Development of power modeling and estimation framework for systems-on-chip.
(ii) Development of optimal allocation and scheduling techniques for energy-efficient mapping of multi-task applications onto multi-processor systems-on-chips.
(iii) Development of energy-scavenging techniques for ultra-low power sensor network platforms.

Team leader: Petru Eles – Linköping University (Sweden)

(i) Development of optimization approaches for energy efficient, time constrained embedded systems.
(ii) Communication synthesis for energy-efficient real-time and fault-tolerant applications implemented on SoC.
(iii) Low power scheduling with voltage scaling and body biasing.

Team Leader: Jan Madsen – Technical University of Denmark (Denmark)

(i) Development low-power asynchronous circuit design, in particular for efficient Network-on-Chip structures.
(ii) Development of multi-objective optimization approaches for exploring the mapping of multitask applications onto multiprocessor System-on-Chip.
(iii) Experience in design and modelling of wireless sensor network platforms.

Team Leader: Lothar Thiele – ETH Zurich (Switzerland)

(i) Combining scheduling methods with energy constraints.
(ii) Extensive experience in performance analysis for real-time embedded systems and interface-based approaches.
(iii) Experience in building and deploying sensor network platforms.

1.2 Affiliated Participants and Roles

Team Leader: Roberto Zafalon – STMicroelectronics (Italy)

(i) Provides an up-to-date view on industrial requirements for low-power system-on-chip platforms.
(ii) Makes available to the activity's participants up-to-date information on the power dissipated in current and up-coming advanced CMOS technologies, as developed by STMicroelectronics.
(iii) Provides information on applications and architecture trends in the market of consume multimedia silicon platforms.

Team Leader: Salvatore Carta – University of Cagliari (Italy)

(i) Cooperates with the cluster's participant on the development of power analysis and modelling techniques for on-chip interconnects (networks-on-chip).
(ii) Provides competences and reference flows for back-end design issues (power issues arising from low-level design – logic design, placement, routing, etc.).

1.3 Starting Date, and Expected Ending Date

Starting date: September 1st, 2004.

Ending date: the activity will span the duration of the project, and continue beyond the end of the project. This is because all current trends indicate that low power design will increasingly become a primary concern and focus of action for researchers, designers and developers working on embedded systems.

Moreover, the integration achieved by this activity is creating the know-how and the skills required for the definition of new research and development initiatives. As a result, all partners are already actively involved in long term funded research programs in low power design, whose end-date is beyond the duration of ARTIST2.

1.4 Baseline

The group of Luca Benini at the **University of Bologna** (UoB) is one of the leading centers in low power design, focusing on system level power management both from the architectural and from the software viewpoint. In this area, the group has produced a large number of contributions on OS-Based-dynamic power management, memory and communication architecture optimization for low power consumption, low power circuit design, battery-driven power management.

One of the baselines contributions of UoB to the project is a complete power modeling infrastructure both for all components of current MPSoC Platforms and for future Network-on-Chip-based platform.

More recently, UoB has focused on power optimization techniques for multi-processor systems on chip, with special emphasis on static and dynamic power management through voltage and frequency setting.

The group of Jan Madsen at the **Technical University of Denmark** (DTU) aims at low-power techniques for wireless sensor networks, and it brings significant experience on low-power asynchronous circuit design, as well as analytic and stochastic modeling of power consumption and battery usage. They have developed exploration methods for mapping applications onto heterogeneous multiprocessor platforms based on a multi-objective optimization framework from ETH Zurich, where one of the objectives is power consumption.

The group of Petru Eles at **Linkoeping University** (LIU) has given important contributions on high-level system modeling of both power and reliability, and on optimization techniques for energy efficient mapping of applications on execution platforms. They have developed approaches for energy efficient implementation of both hard and soft real-time systems.

The group of Lothar Thiele at **ETH Zurich** (ETHZ) has a long standing experience in the area of sensor networks. In particular, a low power platform including hardware, operating system, middleware and various applications has been developed (BTnode). It is used by many research groups worldwide. Together with the experience in real-time systems and scheduling, this is the basis for the joint effort in the design of low power massively distributed systems.

1.5 Problem Tackled in Year 3

Adaptive Power Management for Energy Harvesting Sensor Nodes

Sensor networks are an increasingly important class of distributed embedded systems. Wireless sensor networks are strategic enablers for a number of “ambient intelligence” applications, such as environmental monitoring, monitoring of body functions, tracking of people and objects. The main focus of the work at UoB and ETHZ is on a new class of energy-harvesting devices which replenish the available energy, e.g. by the use of solar cells. The main objective tackled here is the development and test of power management algorithms which adapt to the available environmental energy. The end goal is to enable sustainable environmentally powered operation for sensor networks. In year 3, our efforts were twofold: On the one hand, we have developed new control algorithms which tune the application parameters according to the time-varying amount of harvested energy. Albeit maximizing the utility of a sensor application is a complex optimization problem, we showed how low-complexity controllers may be obtained anyhow. On the other hand, a miniaturized photovoltaic energy harvester has been designed and prototyped which performs automatic power point tracking at a minimum energy cost.

Energy-aware routing in Wireless Sensor Networks

DTU have addressed the problem of energy optimized routing in wireless sensor networks which have energy harvesting capabilities. The have proposed a adaptive energy-aware routing algorithm which is able to route around nodes with low energy, hence allowing these nodes to regain energy before they are included in routing traffic in the network. Intitial results based on simulation have shown that it is possible to gain energy and hence, extend the lifetime of the network, even though the energy-aware routing algorithm uses more energy to analyse and adapt the network traffic than the classical non-energy-aware routing algorithms.

Energy optimization of fault-tolerant embedded systems

DTU and LiU have captured the effect of voltage scaling on system reliability and have shown that if the voltage is lowered to reduce energy consumption, the reliability is significantly reduced. To address this problem, they have proposed a CLP-based approach (Constraint Logic Programming) that takes reliability into account when performing scheduling and voltage scaling. Experimental results have shown, that the CLP-based strategy is able to produce energy-efficient implementations which are schedulable and fault tolerant. By carefully deciding on the start times and voltages of processes they have shown that it is possible to eliminate the negative impact of voltage scaling on reliability without a significant loss of energy savings.

Power optimization and analysis for Nanometer technologies

UoB has addressed the problem of power consumption with Nanometer technologies. Leakage power is a major concern in sub-90nm CMOS technology and numerous design techniques have been proposed to reduce standby leakage in digital circuits. Out of this rich set of solutions, power-gating or sleep transistor insertion has proven to be a very effective approach to reduce standby leakage, while keeping high speed in the active mode. Another interesting solution is based on the usage of multi-threshold libraries and multiple power supplies. While the first approach is considered to be more effective, it has significant impact on design automation flows. The second approaches is less disruptive in terms of design automation, but it often leads reduced savings because it is applied very late in the design process (i.e. during low-level logic synthesis).

Power optimization via system-level resource allocation and scheduling

LiU and UoB have continued to the cooperation on this topic, focusing on energy optimization, and more specifically on the optimization of energy-efficient time constrained multiprocessor

systems. We tackled optimization problems in system level design for variable voltage/frequency MPSoCs, aiming at minimizing power consumption.

The major challenge that we faced was to develop a cooperative solving framework where the allocation, the scheduling and the discrete voltage/frequency selection problem models can be suitably accommodated and solving algorithms integrated.

1.6 Comments From Year 2 Review

1.6.1 Reviewers' Comments

No particular comments related to this activity have been given.

2. Summary of Activity Progress

2.1 Previous Work in Year 1

University of Bologna has focused on interconnect optimisation techniques for low power. Several schemes have been developed to instantiate application (platform) specific interconnect architectures for minimum energy consumption. An algorithm for automatic instantiation of multi-hop busses which includes topology generation and bus frequency assignment has been developed in collaboration with Penn State University.

A research effort on energy aware mapping of multi-task applications on multi-processor SoC execution platforms has been also started. The approach is based on variable-voltage processors where execution speed and voltage supply can be independently adapted to the processor's workload. The first result of this effort has been a design space exploration technique that automatically finds pareto points in the power vs. throughput design space. The technique has been tested on streaming-like signal processing applications. This work is conducted in cooperation with Linköping University.

Additionally several extensions to the power modelling infrastructure in the MPARM virtual platform simulators have been developed, including the model for variable frequency and variable voltage cores, as well as a prototype model for estimating the power consumption of IOs and external memories (this work has been performed in cooperation with associate partner STMicroelectronics).

Linköping University has developed a technique for static routing on NoC, with guaranteed delays and arrival probabilities in the presence of transient faults. The approach is based on schedulability analysis of tasks and messages with priority based arbitration. For fault-tolerance, a combination of spatial and temporal redundancy is considered. Reduced communication energy is one of the goals. More recently the the analysis of the worst-case buffer space needed has been performed. Based on this analysis, it is possible to develop an approach to buffer space minimization in the context described above.

Linköping University's efforts are also aiming at a more accurate modelling of actual communication and memory techniques used in MP SoC. Such an accurate modelling is needed in order for a system level analysis and optimization to produce useful results. Thus, work is concentrating on: (1) Capturing the background communication due to cache misses in system level models. (2) Capturing the bus load due to system-wide synchronization. Once this modelling issues are solved, different optimization techniques can be used for e.g. task mapping and scheduling, as well as voltage selection. Results can be validated using accurate and fast simulation in the environment developed at Bologna. Another issue which is currently addressed is that of efficient optimization techniques based on advanced constraint solving and mathematical programming techniques. This work is performed in cooperation with the group at University of Bologna.

Technical University of Denmark has started the development of a generic sensor network platform (Hogthrob project) which allows to trade-off hardware and software implementations of the various components of the platform. So far, the focus has been on: (1) Processor design: Low-power design techniques have been investigated, included low-power synthesis (e.g., clock-gating), power modes and de-synchronizing in the context of the OpenCores AVR core (using Synopsys) (2) Power modelling: Simulation-based power modelling and estimation techniques have been investigated. This involves analytic and stochastic modelling of batteries and investigation into the macromodeling of various hardware components.

Based on the prototype sensor network platform developed within Hogthrob, various test bench programs have been run on an AVR core synthesized on the FPGA and a number of physical measurements have been conducted. Finally, DTU has started investigating how the power modelling and the sensor network modelling can be captured within the multiprocessor simulation environment, ARTS, developed at DTU.

2.2 Previous Work in Year 2

Power modelling for complex SoC platforms

The activity has focused on extending system-level energy analysis to highly integrated MPSoC platforms with segmented bus architectures, where the efficiency of bridges and protocol/frequency/size converters comes into play to determine the performance of the system interconnect. We leveraged a close cooperation with associate member STMicroelectronics which provided the models, traffic generators, system specifications and performance requirements. Platforms based on the on-chip communication protocols STBus, AMBA AHB, AMBA AXI-have been modeled and simulated at a very high level of accuracy (cycle-accuracy and bus-signal-accuracy), and compared with mixed AHB/AXI platforms

The original MPARM platform allowed the modelling and simulation of single-node communication architectures (as depicted in Fig.1a). The platform was enhanced with the possibility to extend the modelling capability to a multi-layer architecture, as illustrated in Fig.1b. The first scenario corresponds to low-end real-life platforms, where AMBA AHB, AMBA AXI or STBus are the architectures of choice to accommodate on-chip communication. The MPARM platform can also instantiate a NoC as the communication fabric, by wrapping the masters and slaves with the proper network interfaces. In general, all cores can be wrapped with the native bus interface. More complex MPSoC platforms adopt the communication architecture depicted in Fig.1b. It is a hierarchical infrastructure, where communication takes place at a first level of the hierarchy in the local AMBA AHB layers, and at a second level with the system-level slaves. The AMBA Multi-Layer specification introduced the notion of the interconnect matrix first, by envisioning point-arbitration at the destination slaves. This solution is quite interesting, since it allows a larger scalability than single-node solutions. Unfortunately, fabrication problems arise when the number of input layers increases a lot, since the implementation of the interconnect matrix is mostly combinational. This gives rise to clock frequency limitations and to layout unpredictability. As the level of integration of MPSoCs increases, the illustrated structures cannot satisfy communication requirements any more.

A further increase in communication scalability is exposed by segmented architectures, where a number of busses are interconnected with each other by means of bridges. In this case, the congestion on each bus is greatly decreased, thus favouring lower bus access times, but the latency of bus transactions can be seriously increased because of the multiple steps needed to reach a slave located on a different bus. Bridge traversal latency can significantly contribute to overall communication latency. Similarly, the use of bridges raises power concerns. The use of bridges helps to relieve the scalability limitations of traditional communication architectures, however the associated cost consists of the design of a complex IP block (the bridge itself) which is far from trivial and which can significantly affect system performance and energy. Many times, bridges do not perform only protocol conversion, but also size and frequency conversion. In fact, cores with homogeneous characteristics (i.e., clock frequency, data and address bus width) are typically grouped in the same node, therefore each “segment” of the global communication architecture turns out to be a domain with distinctive features. This obviously increases the bridging cost, since up/down size conversions or frequency conversions all take clock cycles to be carried out.

Another issue concerned the porting of traffic generators in order to make the simulation of complex systems in reasonable time possible. Moreover, this allowed overcoming confidentiality problems related to the intellectual property of communicating actors. STMicroelectronics made available its traffic generators for audio and video IP blocks, allowing us to reproduce on the MPSIM environment the traffic patterns of real-life set-top-box platforms with a high level of accuracy.

Another effect of the joint work on traffic generators between Technical university of Denmark and University of Bologna was the development of the necessary infra-structure to co-simulate modules of the abstract system-level MPSoC ARTS frameworks (DTU) with modules available in the cycle-true MPARM framework. The motivation of the work is to investigate MPSoC instances at mixed-levels of abstraction. A simple system where two ARTS IP cores were connected through a MPARM AMBA-AHB bus was successfully implemented and co-simulated.

Finally, a significant modelling effort was required also for the memory controller. In fact, MPSIM has traditionally simulated MPSoC systems with on-chip memories only; therefore we needed to model real-life memory controllers for I/O. We got the LMI specification from STMicroelectronics, and developed a SystemC model which was accurately (cycle-by-cycle) validated against the behaviour of the real LMI. Such powerful model allows us to interface our MPSoC with SDR and DDR SDRAMs, and more interestingly to model I/O access latency of real systems. Finally, we retain the capability to model an on-chip shared memory in place of the off-chip SDRAM, thus being able to differentiate system performance and power in presence of a slow off-chip memory vs. a fast on-chip memory. Optimizations for access to the off-chip memory can also be analyzed with this platform.

Power optimization via system-level resource allocation and scheduling

In this activity, the focus is on addressing resource allocation problems in Multi-Processor Systems-on-Chip (MPSoCs). An important instance of this problem is when have to allocate and schedule a given task graph (representing a functional abstraction of a multi-task application) on a target multi core platform while choosing the frequency (and voltage) at which each task will be executed. Since hardware platforms and applications are extremely complex, it becomes thus important not only to measure the optimizer efficiency as done in general in the optimization area, but also to verify if the optimization model is accurate through a validation step performed via simulation on a virtual platform.

Allocation, scheduling and discrete voltage selection problem for variable voltage/ frequency MPSoCs, minimizing the system energy dissipation and the overhead for frequency switching, are clearly NP-hard problems. Only incomplete approaches have been proposed to solve these problems in the system design community. In this activity we have investigated a hybrid methodology based both on Constraint Programming (CP) and Integer Programming (IP) that splits the overall problem in two subproblems, the first being the allocation of tasks to processors and frequencies to tasks and the second being the scheduling. Our methodology derives static allocation, scheduling and frequency setting; therefore it targets applications with design-time predictable behaviour.

In order to solve the problem to optimality without incurring accuracy limitations, we applied the concept behind the *logic-based Benders decomposition technique* to this new application problem. Bender decomposition can be summarized as follows. A complex optimization problem is decomposed in two parts: the first, called Master Problem, is the allocation of processors and frequencies to tasks and the second, called Subproblem, is the scheduling of tasks given the static allocation and frequency assignments provided by the master. The master problem is tackled by an Integer Programming solver while the subproblem through a Constraint Programming solver. The two solvers interact via generation of no-goods

(constraints on acceptable solutions for the CP solver) and cutting planes (constraints on acceptable values of the integer variables for the IP solver) generation. The solution of the master is passed to the subproblem in an iterative procedure that is proved to converge to the optimal solution. This work has been performed in cooperation by the University of Bologna and Lonköping University.

The methodology has been tested on a variety of realistic instances. In addition, we test the accuracy of the solutions provided by the optimizer simulating them on an MPSoC virtual platform. In particular, we have used two demonstrators (GSM and JPEG) to prove the applicability of the developed methodology to real-life embedded applications scenarios.

In a parallel, but strongly related activity, we have also addressed the specific problems of soft real-time systems. In this case, certain tasks are allowed to miss their deadlines. This however, negatively affects the delivered QoS. The goal is to maximize the QoS with a limited energy budget or to achieve a certain level of QoS with as low energy consumption as possible. Linköping University has developed heuristics which determine the system schedule and voltage levels of tasks in such a system.

Finally, DTU has experimented with the use of meta-heuristics to solve the 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. Our approach allows for mapping onto a fixed platform or onto a flexible platform where architectural changes are explored during the mapping. The approach uses multi-objective evolutionary algorithms and is based on the PISA framework for multi-objective optimization developed at ETH Zurich. We demonstrate the approach through an exploration of a smart phone, where five task graphs with a total of 530 tasks after hyper period extension are mapped onto a multiprocessor platform. The results show four non-inferior solutions out of 10.000 explored solutions, which tradeoffs the various objectives.

Scheduling based energy optimization for energy-scavenging wireless sensor networks

Wireless sensor networks – consisting of numerous tiny sensors that are unobtrusively embedded in their environment – have been the subject of intensive research. As for many other battery-operated embedded systems, a sensor's operating time is a crucial design parameter. As electronic systems continue to shrink, however, less energy is storable on-board. Research continues to develop higher energy-density batteries and supercapacitors, but the amount of energy available still severely limits the system's lifespan. As a result, size and weight of most existing sensor nodes are largely dominated by their batteries.

On the other hand, one of the main advantages of wireless sensor networks is their independence of pre-established infrastructure. That is, in most common scenarios, recharging or replacing nodes' batteries is not practical due to (a) inaccessibility and/or (b) sheer number of the sensor nodes. In order for sensor networks to become a ubiquitous part of our environment, alternative power sources should be employed. Therefore, environmental energy harvesting is deemed a promising approach: If nodes are equipped with energy transducers like e.g. solar cells, the generated energy may increase the autonomy of the nodes significantly. Several technologies have been discussed how, e.g., solar, thermal, kinetic or vibrational energy may be extracted from a node's physical environment. Moreover, several prototypes have been presented which demonstrate both feasibility and usefulness of sensors nodes which are powered by solar or vibrational energy.

The focus of this activity is on sensor nodes with energy-scavenging features. In general our results apply for all kind of energy harvesting systems which must schedule processes under deadline constraints. For these systems, new scheduling disciplines must be tailored to the

energy-driven nature of the problem. This insight originates from the fact, that energy – contrary to the computation resource "time" – is storable. As a consequence, every time we withdraw energy from the battery to execute a task, we change the state of our scheduling system. That is, after having scheduled a first task the next task will encounter a lower energy level in the system which in turn will affect its own execution. This is not the case in conventional real-time scheduling where time just elapses either used or unused.

The main developments obtained in this activity can be summarized as follows

(a) We studied an energy-driven scheduling scenario for a system whose energy storage is recharged by an environmental source. For this scenario, we developed an optimal online algorithm that dynamically assigns power to arriving tasks. These algorithms are “energy-clairvoyant”, i.e., scheduling decisions are driven by the knowledge of the future incoming energy.

(b) We developed an admittance test that decides, whether a set of tasks can be scheduled with the energy produced by the harvesting unit, taking into account both energy and time constraints. For this purpose, we introduced the concept of energy variability characterization curves (EVCC).

(c) In addition, a comparison to earliest-deadline first (EDF) by means of simulation, demonstrated that significant capacity savings can be achieved by our approach, when compared to the classical EDF algorithm.

2.3 Current Results

2.3.1 Technical Achievements

Power optimization and analysis for Nanometer technologies (UoB)

Leakage power is a major concern in sub-90nm CMOS technology and numerous design techniques have been proposed to reduce standby leakage in digital circuits. Out of this rich set of solutions, power-gating or sleep transistor insertion has proven to be a very effective approach to reduce standby leakage, while keeping high speed in the active mode. Another interesting solution is based on the usage of multi-threshold libraries and multiple power supplies. While the first approach is considered to be more effective, it has significant impact on design automation flows. The second approaches is less disruptive in terms of design automation, but it often leads reduced savings because it is applied very late in the design process (i.e. during low-level logic synthesis).

The work performed by UoB aimed at addressing the above mentioned limitations of power-gating and multi-threshold/voltage synthesis. The work has been performed in cooperation with associate partner STMicrolectronics, Bulldast and Politecnico di Torino, in close cooperation with the 6th FP integrated project ICT-CLEAN. The contributions have been in two main directions:

a. Developing design automation techniques for automatic power gating insertion. The following issues have been tackled: (i) clustering of non-timing critical cells in a physical design aware fashion to enable partial power gating with low area overhead. The key idea is to prevent power gating of timing critical cells that would decrease the speed of the circuit. (ii) sleep transistor sizing and insertion. The contribution is to automatically size sleep transistor cells for a specified speed within area constraints. Results on 65nm libraries have demonstrated that significant power savings can be achieved even

for circuits that are not amenable to full-power gating because of tight timing constraints (up to 80% of the full power savings potential without any speed degradation has been obtained).

b. Development of high-level (register-transfer level) strategies for automatically exploring the speed vs. power trade-off for arithmetic units in datapath circuits. This work enables a leakage-power aware arithmetic unit instantiation strategy during the early phases of the synthesis process. Given a timing constraint and an average activity information, the new algorithm is capable of selecting the most energy efficient functional unit (e.g. carry lookahead vs. carry propagate adder) in an automatic fashion. Experiments have demonstrated that significant improvements can be achieved in energy efficiency compared with the traditional strategy that optimizes for leakage power only during the last steps of logic synthesis. In average 20% extra leakage savings have been obtained in a 65nm library with respect to standard logic synthesis techniques.

These work represent the first steps in developing new methods and tools for designing embedded systems in leakage-dominated technologies, a topic that will be of growing interest for the embedded system design community. This work has particular relevance for the ARTIST2 activity **Power optimization via system-level resource allocation and scheduling,** as the high-level allocation and scheduling decisions will be increasingly influenced by leakage power, and by the availability of low-leakage shutdown states.

The main WEB reference to the work performed in this area is the web site of the CLEAN project: http://clean.offis.de The following papers detail the technical achievements obtained with the participation of UoB.

Power optimization via system-level resource allocation and scheduling (LiU, UoB)

Linkoeping and Bologna have continued to the cooperation on this topic, focusing on energy optimization, and more specifically on the optimization of energy-efficient time constrained multiprocessor systems. We tackled optimization problems in system level design for variable voltage/frequency MPSoCs, aiming at minimizing power consumption.

The major challenge that we faced was to develop a cooperative solving framework where the allocation, the scheduling and the discrete voltage/frequency selection problem models can be suitably accommodated and solving algorithms integrated.

By leveraging the principle of logic-based Benders Decomposition we created with an iterative two-step mapping framework that is proved to converge to the optimal solution. Our methodology derives static allocation, scheduling and frequency settings, therefore targets applications with design-time predictable behaviour. Signal processing and multimedia applications employing pipelining as workload allocation policy are the most common example of such applications. Therefore our optimizer was tuned for energy-efficient mapping of pipelined task graphs on MPSoCs. Finally, we used two demonstrators to prove the applicability of the developed methodology to real-life scenarios.

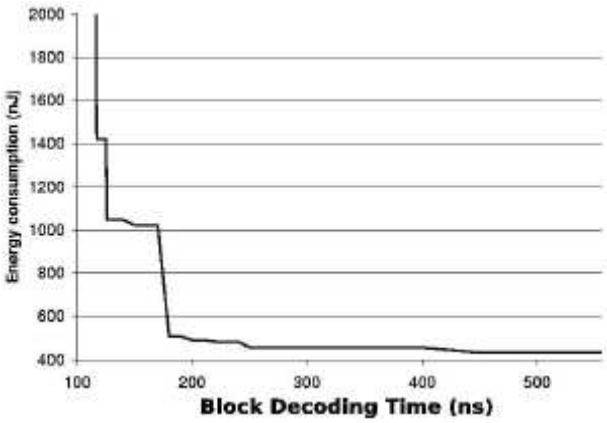

The above figure shows the Pareto-optimal set of solutions, spanning the trade-off between decoding time and Energy consumption for a GSM decoding application running on a multi-core ARM platform. The solution points in the graph where computed by running the optimizer with different block decoding time constraints. As it can bee seen, there is a wide range of feasible operating points that can be automatically computed by the optimizer. It is important to emphasise that both energy and block decoding times have been validate with cycle accurate simulation and power estimation on a virtual platform. Hence, the accuracy of the optimization model has been carefully assessed.

Adaptive Power Management for Energy Harvesting Sensor Nodes (UoB and ETHZ)

Energy is a primary constraint in the design of sensor networks. This fundamental energy constraint further limits everything from data sensing rates and link bandwidth, to node size and weight. For example, it has become evident that energy storage devices largely dominate the form factor of existing sensor nodes and prevent their miniaturization. Moreover, the limitation of the energy supply has constantly impeded the progress of sensor networks towards large scales and true autonomous operation. For instance, long-term monitoring of the environment is one of the sensor network visions. For this application, finite energy reservoirs like batteries render the deployment and maintenance of large scale sensor networks extremely cumbersome.

The focus of this activity is on sensor nodes with energy-scavenging features. Several technologies have been proposed in the past years how, e.g., solar, thermal, kinetic or vibrational energy may be extracted from a node‘s physical environment. Solar energy is certainly one of the most obvious and promising energy sources. Clearly, one may just use solar energy to recharge a primary energy source, like e.g. a battery. Like that, the point in time when the system runs out of energy is simply postponed. If, however, one strives for perpetual operation, common power management techniques have to be reconceived. In addition to perform classical power saving techniques, the sensor node has to adapt to the stochastic nature of the solar energy. Goal of this adaptation is to maximize the utility of the application in a long-term perspective. The resulting mode of operation is sometimes also called *energy neutral* operation: The performance of the application is not predetermined a priori, but adjusted in a best effort manner during runtime and ultimately dictated by the power source.

The main developments achieved in this activity can be summarized as follows:

(a) We present feedback controllers which adapt task rates such that maximal utility is obtained while respecting the time-varying amount of harvested energy. Instead of solving the

optimization problem on-line which may be prohibitively complex in terms of running time and power consumption, we propose the use of multiparametric, optimal on-line controllers.

(b) We present a hierarchical control design which overcomes several drawbacks of previously proposed designs: First, the computation overhead and storage demand of the online controller is reduced significantly. In a case study, the achieved reductions compared to a single controller amount 91% and 83%, respectively. Second, by designing the controller for worst-case situations, depletion of the energy storage is avoided and robustness of the overall system is increased.

(c) For some applications of practical concern, the optimal multiparametric solutions may grow to complex for constraint systems like sensor nodes. For those applications, we propose a novel algorithm for approximate multiparametric linear programming. We demonstrate, that the online complexity of the generated control laws is highly reduced compared to an optimal solution in terms of computation overhead and storage demand. For an example investigated, we found improvements of 98% and 92%, respectively. Furthermore, simulations show that the the performance of the found control laws is comparable with the one achieved by complex, optimal control laws.

(d) We propose a practical technique for the efficient implementation of the controller and demonstrate the practical relevance of our approach by measurements of the controller running on a real sensor node.

(e) Beside this theoretical framework, a novel analog circuit for solar energy harvesting has been presented. Here, a key design challenge is how to optimize the efficiency of the solar energy collection under non stationary light conditions. The proposed scavenger exploits miniaturized photovoltaic modules to perform automatic maximum power point tracking.

Energy-aware routing in Wireless Sensor Networks (DTU)

The Technical University of Denmark (DTU) have studied low power aspects of wireless sensor networks, in particular for networks where the nodes are able to harvest energy from the environment and hence being able to recharge their batteries. DTU have developed an adaptable energy-aware routing algorithm based on directed diffusion which is able to avoid routing through nodes with low energy. This strategy allows low energy nodes to recover through energy harvesting. When such a node reaches a certain energy level, it may again be considered for routing data. The additional energy usage introduced by the more advanced routing algorithm is spent for a smart distribution of traffic which takes into account dynamic changes of node status and energy harvesting for the benefit of the network performance and lifetime. The routing algorithm has been implemented and simulated in a setup which uses a solar panel as the energy harvesting unit, and simulates the environment by varying the amount of energy harvested over day and night as well as by the introduction of additional factors like natural shadows which are influencing the amount of energy being harvested. Such kind of fluctuations are noticed by the algorithm and used to arrange traffic and energy consumption to avoid node areas of lower energy.

Communication between nodes in wireless sensor networks is susceptible to transmission errors caused by low signal strength or interference. These errors manifest themselves as lost or corrupt packets. This often leads to retransmission, which in turn results in increased power consumption reducing node and network lifetime. DTU have implemented and evaluated a convolution code FEC with Viterbi decoding on Mica2 nodes to explore the possibility of extending the lifetime of a degrading wireless sensor network. Results from these experiments [DOM06] suggest that the approach could be used in a wireless sensor network when increasing distance and channel noise due to node dropout, degrade the network.

Energy optimization of fault-tolerant embedded systems (LiU, DTU)

Addressing simultaneously energy and reliability is especially challenging because lowering the voltage to reduce energy consumption has been shown to exponentially increase the number of transient faults. The main reason for such an increase is that, for lower voltages, even very low energy particles are likely to create a critical charge that leads to a transient fault. However, this aspect has received very limited attention. Moreover, time-redundancy based fault-tolerance techniques, such as re-execution, and voltage scaling-based low-power techniques are both relying on the use of processor idle-time. In addition, such competing requirements have to be met within a given development and manufacturing cost and time-frame. Therefore, the task of designing such embedded systems is becoming not only increasingly important, but also increasingly difficult.

Linkoeping University (LiU) has proposed a technique to combine re-execution and active replication in an optimized implementation that leads to a schedulable fault-tolerant application without increasing the amount of employed resources. They have also addressed transparency/performance trade-offs during the synthesis of fault-tolerant schedules. Starting from their research, the Technical University of Denmark (DTU) and LiU have jointly considered a very different trade-off, namely, energy versus reliability.

They have 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 [Pop07a, Pou07b]. This research has considered heterogeneous distributed time-triggered systems, where both processes and messages are statically scheduled. The transient faults are tolerated through process re-execution by switching to pre-determined contingency schedules.

2.3.2 Individual Publications Resulting from these Achievements

UoB

- Sathanur, A. Pullini, L. Benini*, A. Macii, E. Macii and M. Poncino “Timing Driven Row-Based Power Gating” International Symposium on Low power Electronics and Design 2007.
- S. Merdardoni, D. Bertozzi, E. Macii, “Power-Optimal RTL Arithmetic Unit Soft-Macro Selection Strategy for Leakage-Sensitive Technologies”. International Symposium on Low Power Electronics and Design 2007.

DTU

- S. U. Donapudi, C. O. Obel, J. Madsen, ”Extending Lifetime of Wireless Sensor Networks using Forward Error Correction”, Norchip 2006.

LiU

- A. Andrei, P. Eles, Z. Peng, M. Schmitz, B. Al-Hashimi, "Voltage Selection for Time-Constrained Multiprocessor Systems on Chip", Chapter in "Designing Embedded Processors: A Low Power Perspective", eds. J. Henkel, S.Parameswaran, Springer 2007.

- A. Andrei, P. Eles, Z. Peng, M. Schmitz, B. Al-Hashimi, "Energy Optimization of Multiprocessor Systems on Chip by Voltage Selection", IEEE Transactions on Very Large Scale Integration Systems, V15, Nr. 3, March 2007.
-

2.3.3 Interaction and Building Excellence between Partners

The reported results have been produced as a consequence of systematic and intensive interaction between the partners in the cluster. This includes both the industrial and the academic partners. Master students, PhD students, and senior research staff have participated in mobilities. Meetings have been organised e.g. at DATE 2007 and during the cluster meeting in Linköping, May 2007. Software platforms have been exchanged between the partners. Several joint publications have been produced

Power optimization and analysis for Nanometer technologies (UoB, ST Microelectronics)

The interaction in this activity has been mostly with associate industrial and academic partners. The associate partners involved provided the necessary link with advanced silicon technology and related information. This is essential background required to build high-level models and a higher level understanding of the nature of the technical problem. A clear understanding of the low-level leakage sources and the implementation of leakage reduction techniques is essential prerequisite for developing higher-level power optimization policies (i.e. allocation and scheduling) that take into account leakage.

Adaptive Power Management for Energy Harvesting Sensor Nodes (UoB, ETH Zurich)

The main focus of the work at UoB and ETHZ is on a new class of energy-harvesting devices which replenish the available energy, e.g. by the use of solar cells. The close cooperation has been performed during several research exchanges and has resulted in common publication.

- Meeting/discussion at the conference “Design, Automation and Test in Europe (DATE07)”, Nice, France, April 16-20, 2007.
- Davide Brunelli from the University of Bologna visits the research group at ETH Zurich from 7th July until 28th of September 2007 for scientific collaboration.
- Master thesis student Steven Laurier from ETH Zurich visits the research group in Bologna for information/technology exchange
- Hardware technology exchange: solar scavenging prototype for sensor nodes shipped from Bologna to Zurich, BTnode developments kits sent vice versa.

Power optimization via system-level resource allocation and scheduling (LiU, UoB)

The interaction in this activity has been between Linkoeping and UoB. Integration from the technical viewpoint has been on the development of the optimization model, and on creating an executable environment that matches accurately the optimization model. Previous experience of Linkoeping on task allocation and scheduling modelling has been essential in this phase. UoB has contributed with its experience in software development and manual mapping of applications on multi-core platforms, as well as its experience in complete optimization algorithms for scheduling and allocation.

- Meeting/discussion at the conference “Design, Automation and Test in Europe (DATE07)”, Nice, France, April 16-20, 2007.

- Simulation technology transfer.
- Several experiments in cooperation and paper writing.

Energy optimization of fault-tolerant embedded systems (LiU, DTU)

The interaction in this activity has been between Linköping and DTU. This is part of a broad cooperation between the two groups in the area of fault tolerant real-time systems. Two common publications have resulted.

- Meeting/discussion at the conference "Design, Automation and Test in Europe (DATE07)", Nice, France, April 16-20, 2007.
- Meeting at Linköping (June 2007); visit of Paul Pop from DTU.
- A PhD student from Linköping has spent several weeks at DTU.

2.3.4 Joint Publications Resulting from these Achievements

- 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.
- C. Moser, L. Thiele, D. Brunelli, and L. Benini, "Adaptive Power Management in Energy Harvesting Systems.", In *Design, Automation and Test in Europe (DATE 07)*, Nice, France, April 16-20, 2007.
- 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.
- 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

2.3.5 Keynotes, Workshops, Tutorials

The following keynote and tutorial have been delivered around the topic of this activity. They have contributed to the dissemination of the research results to the broad community, both in industry and academia, and also beyond the borders of Europe.

Keynote: "Chips i alt" (eng. "Chips everywhere")
Tåstrup, Denmark – May 30, 2007

Speaker: Jan Madsen

The Danish Academy of Technical Sciences held a one-day seminar with the title "Chips everywhere". Jan Madsen gave a keynote speech on the system challenges of designing wireless sensor networks, particular emphasizing the challenges of making these systems energy-aware in order to extend their life-times. The seminar was attended by approximately 50 people, of which most were from companies with strong interests in embedded systems.

Tutorial: "Low Power CMOS Design: The Fabrics: Research Front-end"
Asian-Pacific Design Automation Conference**,** *Yokohama, Japan, January 23 2007.*

The tutorial coverer research front-ends of low power CMOS design, including (1) process and device level, (2) circuit level, (3) EDA level, and (4) system level. The focus of the presentation given by Luca Benini was on system-level power optimization

http://www.aspdac.com/aspdac2007/tutorial/index.html

3. Future Work and Evolution

3.1 Problem to be Tackled over the next 12 months (Sept 2007 – Aug 2008)

Power optimization via system-level resource allocation and scheduling

In this area, future work will be directed toward generalizing the model to take into account also deep submicron effects, such as leakage power high failure rates and variations. Moreover, a generalization on the application models will be pursued: more general types of task graphs will also be targeted by the optimal allocator and scheduler. Finally a comparison between exact and heuristic methods will be performed. In summary the problems to be tackled are:

- Power optimization for nanometer platforms. Models should be extended to deal with limited number of operating points, leakage power, and 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

Adaptive Power Management for Energy Harvesting Sensor Nodes

By implementing the algorithms developed at ETH Zurich on the solar scavenger prototype developed at the University of Bologna, we plan 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 – we plan to demonstrate the usefulness of our theoretical results in a practical application, possibly in an outdoor setting

Energy optimization of fault-tolerant embedded systems (LiU, DTU)

Further collaboration of DTU and LiU will focus on energy and reliability of embedded systems. Once a scheduling technique for low-power reliable applications is available, the impact of different redundancy techniques can be investigated in detail. They will first determine the effect of re-execution and replication on the energy consumption. Then they will propose reliability analysis techniques for redundancy. They will investigate redundancy assignment optimization approaches to increase the reliability at a reduced energy cost.

Temperature-aware Voltage Selection for Energy Optimization

The high power densities achieved in current SoCs do not only result in huge energy consumption but also lead to increased chip temperatures. High temperatures can impact reliability as well as cooling and package cost. One aspect, of particular interest, is that growing temperature leads to an increase in leakage power and, consequently, energy, which, again, produces higher temperatures. Research at LiU will investigate techniques for

temperature aware energy minimization by DVS, considering both supply voltage selection and adaptive body biasing.

3.2 Current and Future Milestones

Power optimization via system-level resource allocation and scheduling

Linkoeping and Bologna will continue cooperation on this topic, and we will explore the possibility of including the evolutionary exploration from DTU. Our main goal for the coming period is to consolidate the results regarding the optimization of energy-efficient time constrained multiprocessor systems. The main directions are the following:

- improve and refine the task-based application models as well as the architecture models, in order to make them as realistic as possible, in the context of current execution platforms and target applications; *significant progress has been done in this direction, as outlined above and in the quoted publication. Validation using cycle accurate simulation has been carried out.* ***Further refinements of the models to include leakage power as well as variations will be continued next year, based on the initial work performed this year.***
- explore more efficient design space exploration approaches based on mathematical programming or heuristics, e.g. evolutionary algorithms: *a number of advanced speedup techniques have been implemented to increase the efficiency of Benders-based logic decomposition, and task graphs with up to hundred tasks can be mapped and scheduled by the optimizer.* ***Comparison with heuristics will be the focus in next year****.*
- extend the approaches to on-line voltage scaling, such that dynamic slack can be exploited; the problem is particularly interesting in the context of multiprocessors; *the run-time infrastructure for supporting dynamic task reclamation has been implemented, but preliminary experiments have shown that little extra saving is possible if execution times are highly deterministic.* ***Extension to more non-deterministic task models will be the focus of the next year****.*
- explore interaction and tradeoffs between energy efficiency and fault tolerance: *Significant progress has been made in this direction. Liu and DTU have 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. Two publications have been produced.* **The impact of different redundancy techniques will be investigated in detail. Redundancy assignment optimization approaches will be explored, to increase the reliability at a reduced energy cost.**

Scheduling based energy optimization for energy-scavenging wireless sensor networks

ETHZ and Bologna will continue to work on low power sensor network design. In particular, the results so far will be extended towards application-level decisions.

- To this end, on-line control strategies need to be developed that change the state of the application depending on the current systems state, e.g. the amount of local data

stored, and on the estimation of the future energy flow. The approach will be based on the experience in UoB on building energy-harvesting nodes and on convex optimization from ETHZ. Both theoretical and practical aspects of the problem will be investigated. The feasibility of implementation of advanced on-line control strategies on tightly constrained sensor network hardware platforms will be explored. This activity will leverage the experience on ETHZ and UoB on the hardware-software design of wireless sensor nodes.

UoB and ETHZ are actively developing a joint prototype sensor node which is powered by solar energy. To this end, a BTnode - originally developed at ETH Zurich - has been transferred to Bologna. Equipped with solar panels and supercapacitors for energy storage, measurements on this prototype will illuminate the practical relevance of our theoretical results.

The milestone can be regarded as achieved in a sense that both a theoretical solution for the control problem as well as an efficient solar energy scavenging circuit have been developed by the groups from ETH Zurich and University of Bologna, respectively. However, the merging of both works and the integration into a autonomously running, solar driven sensor node remains to be completed..

Extensive measurements of the scavenger prototype circuit will form the basis for an accurate, application-specific model of the physical hardware. This model will then be used on the algorithmic layer to optimize the overall system performance. The practical relevance of the developed control laws will be demonstrated on a jointly developed sensor network platform.

Power optimization via system-level resource allocation and scheduling

The work on multi-objective design space exploration environment developed at DTU and based on the PISA environment for multi-objective optimization from the group of Lothar Thiele, ETH Zurich, was completed with an invited talk at DIPES 2006, October 2006, Braga, Portugal and a publication.

3.3 Indicators for Integration

Consistent progress has been reported with respect to the integration indicators. Active and productive cooperation between has been further developed:

- The Linköping group has integrated the Bologna simulation platform for both research and education purposes. Several applications have been implemented and research results have been successfully validated. The platform is also used as part of a master-level course for design space exploration projects. The two groups have successfully cooperated with good synergy effect, resulted in research results on system level optimization, new ideas in system modelling, and common publication. This work has required frequent communications between graduate students and a visit of Linkoeping researchers to Bologna. Additional integration effort in year 3 has been focused on developing an accurate and effective optimization model for allocation and scheduling. This has required meetings and interaction through teleconferencing and email. PETRU: anything to add here?
- DTU and UoB have worked closely on the integration of traffic generators in the MPARM simulation platform, and they also achieved the integration of the high-level SystemC models of task execution developed in the ARTS framework, into the MPARM cycle accurate simulation platform. This work has required several weeks of visits of a

graduate student form DTU to the University of Bologna. The cooperation has resulted in joint publications.

- ETHZ and Bologna have worked in close cooperation on energy conservation and energy scavenging for wireless sensor networks. This cooperation has been based on a 6-months stay of a graduate student from Bologna in ETHZ. Moreover, there has been a visit of the primary investigator from bologna to ETHZ. Several joint publications have been produced.
- Linköping and DTU have a strong research collaboration on design and optimization of fault-tolerant embedded systems with the aim to consider efficient inclusion of fault tolerance without increasing energy consumption. Two joint publications have already been published. A PhD student from Linköping has visited DTU for several weeks.
- DTU / INRIA: Joint activities with the Real-Time Components Cluster are currently being explored. Alain Girault from INRIA, France, visited DTU August 23-24. to discuss common issues in fault-tolerant embedded systems.

From the technical point of view, several new problems have been identified, and will be jointly researched by the partners. The research approach strongly leverages synergies between the partners, by integrating different levels of system abstraction (from scheduling via operating systems to system design). The successful technical cooperation is demonstrated by several joint publications.

3.4 Main Funding

Bologna University:

- STMicroelectronics industrial grant, Freescale semiconductor industrial grant

Linkoeping University:

- Swedish Foundation for Strategic Research (SSF)

Technical University of Denmark

- Hogthrob, The Technical Research Council of Denmark (STVF)

ETH Zurich

- Siemens Building Technology Industrial Grant, National Project MICS (Mobile Information and Communication Systems).

4. Internal Reviewers for this Deliverable

Prof. Zebo Peng, department of Computer and Information Science, Linköping University