# Hard Real-Time Embedded Multiprocessor Platform - RTEMP

## 1 Introduction and Objectives

Hard real-time embedded systems are computer based systems where software provides intelligent functionality, and where a failure to respond in time may have fatal consequences. Examples are medical devices, factory automation systems, motor control systems, fly-by-wire systems in airplanes, and brake systems in cars. The design of such systems is driven by the ability to *analyze* and *guarantee* worst-case execution time (WCET) of the system. This calls for processors, compilers etc. which are different from those used in general purpose computing.

In all fields of computing we are witnessing a shift **from single-processor systems towards multiprocessor systems**. This shift is driven by several factors including: (1) The evolution of chip technology that allow tenths to hundreds of processors to be integrated on the same chip. (2) An increased consensus that the use of many smaller and simpler processors is more efficient than the use of one or a few more sophisticated and powerful processors – it represents a better compromise between hardware cost, energy consumption, and compute performance.

The hard real-time industry is naturally conservative. Researchers are just beginning to address the use of multiprocessor platforms and this has many interesting and far-reaching perspectives for the design of real-time systems. Taken to the extreme, **every task can execute on its own processor**. This may greatly **simplify WCET analysis, reduce the response time** of tasks, and allow a more modular design process. The overall objectives of the project are: (1) to develop a **time-predictable multiprocessor platform** for real-time embedded systems that is specifically optimized for implementation in field-programmable gate array (FPGA) technology, and (2) to evaluate this platform using a **real-world application** provided by Danfoss Power Electronics A/S who is a partner in the project. This involves a considerable hardware development effort. We request funding for 1 PhD and 1 Postdoc for 36 months, and one visiting researcher for three months. The project goals are ambitious but realistic – the project represents a coordinated continuation of past research in the two fields of time-predictable processor design [24, 26] and on-chip networks [3, 5]. Furthermore, the project will benefit from synergy with the EU FP7 project T-CREST [34], where DTU has the technical leadership.

The expected results from the project are a multiprocessor platform that **improves WCET** of an existing Danfoss system by **at least a factor of 4**; and achieves this **using less hardware resources and less energy**. An additional success criteria is industrial take-up of the developed technology. The results are of interest for a wide range of industrial companies (including SMEs) engaged in development of embedded systems. To the extent possible, the project outcome will be made available under an industry friendly open-source license.

## 2 Background and State-of-the-Art

**Real-Time Systems = Time-predictable Systems** In real-time embedded systems all components (processor, on-chip network, and memory) must be time predictable. This is in contrast to general-purpose processing where the focus is on raw average-case speed. Many of the features used to achieve this make the execution time highly unpredictable. Examples of such features are: speculation/prediction, cache-replacement strategies, and non-deterministic access to shared resources.

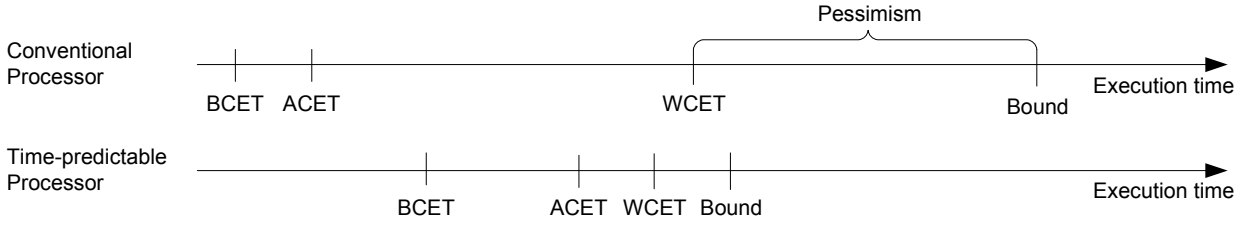


Figure 1: Illustration of the different execution times of a task on different processors.

Figure 1 illustrates the fundamental problem which we address. It shows a single processor, but the challenges apply for multiprocessors as well. For a given program, a conventional processor has a low average-case execution time (ACET) and an even lower best-case execution time (BCET). But the worst-case execution time (WCET) is high. And even worse, it is very difficult to provide an upper bound on the execution time, the estimated WCET bound is often overly pessimistic. When designing a real-time embedded system, it is the WCET Bound and not the ACET, which is of importance. This drives the **design of hardware platforms** for real-time systems down **a completely different path** than the traditional "greed for raw speed" mind-set, and this is the essence and challenge in the proposed project. In [12] Edwards and Lee argue: "It is time for a new era of processors whose temporal behavior is as easily controlled as their logical function". They propose a design called precision timed machine (PRET) [18] for repeatable timing. Problematic architectural features and solutions are presented in the literature [15, 26, 37, 35, 8].

In a multiprocessor platform the network-on-chip as well as a possible external memory are shared resources and therefore potential sources of unpredictable timing behavior. The focus on these new challenges in real-time systems is a distinguishing feature of the project.

**Multi-core Platforms and On-chip Networks** Today, hardware platforms for embedded systems are typically designed as so-called multi-core systems integrating a number of (predesigned) general-purpose processors, digital signal processors, hardware accelerators, memory modules, and IO-units all connected by some form of interconnect fabric. Examples of such systems are cell-phones, set-top boxes for TV, digital cameras, and hearing aids [7, 30]. For the interconnect we are seing a shift

from circuit-switched time-shared busses to packet switched on-chip networks [9, 6]. Since 2006 an ACM/IEEE conference (NOCS) has been devoted to this subject.

Fundamentally, the on-chip network is a shared resource, and this may negatively impact WCET analysis. Just as is the case for processors, most on-chip networks are designed with a focus on average-case speed, and these networks are completely unsuited for use in real-time embedded systems. A time-predictable on-chip network has basically two options: non-blocking routers with rate control (e.g. Mango [1]), and circuit switching (SoCBUS [36]), possibly implemented using time-slicing (Aethereal [13]). In this project, where we are targeting FPGA technology, our aim is the simplest possible design. We have explored initial ideas for a time-predictable on-chip network targeting FPGA-implementation [27] (to be be presented at NOCS'12 in May 2012).

**FPGA Technology** The cost of designing and fabricating an integrated circuit has increased. Today very few system providers can afford to design their own application specific integrated circuits (ASIC). The very high engineering costs simply cannot be amortized over the production volume of the system. Application domain specific standard chips are therefore emerging as one solution to this problem (e.g., in the cell-phone area). Another solution is to use a field-programmable gate array (FPGA) chips where the wiring and the basic logic functions can be configured by the user.

Due to the evolution of chip technology, FPGAs have now become so big that they can be used to implement complete multiprocessor single-chip systems [20], and this development will continue. Therefore, **FPGAs are increasingly being used in products** – a trend particularly **visible in the Danish electronics industry**, which is dominated by small and medium sized enterprises. (10 years ago many Danish companies designed ASICs. Today it is more or less limited to the hearing aid companies, because of their very demanding requirements with respect to size and energy consumption.) The downside of FPGAs compared to ASICs is higher power consumption, somewhat lower performance, and a higher price for high volume products.

Large FPGAs face similar issues with clocking and power consumption as ASICs. One outcome of the project is insight into the development process for complex designs in FPGA technology.

**Previous Research** Our previous results in the field of processor design include development of a highly time-predictable Java processor [24], contributions to profiles for predictable Java [16], and development of analysis tools [28, 14]. Our previous results in the field of on-chip networks – achieved through 5 PhD-projects and 10 years of research at DTU Informatics – include the development of a network that supports both best-effort and guaranteed-service traffic [1, 2, 3], work on simpler and more hardware efficient networks [31, 5, 32], work on programming models and the processor-to-network interface [4, 21], and tools supporting the design of network-on-chip topologies [19, 33].

**Related Projects** We are the technical lead in the EU FP7 project T-CREST [34]. This project addresses a wide range of challenges from analysis tools to hardware implementation. The proposed project will benefit from the T-CREST project, and the proposed project supplements it by focusing on FPGA technology and strong involvement from Danish industry. We are closely following a number of related research efforts in the field of embedded systems, including the Artemis project RECOMP [22] and the DaNES network of excellence [10] supported by the High-Technology Foundation. These activities are focused on software issues and high level models.

## 3 Research Plan

Most successful engineering research is characterized by involving: (i) concepts and theories, (ii) technology, and (iii) applications – the latter to help identify requirements in the early phases of the project and to provide a basis for evaluation in the late phases of the project. For this reason the project involves an industrial partner - Danfoss Power Electronics A/S - who will provide motor and process control use cases (typical application within industrial or food processing area). One or more of these will actually be implemented on the multi-core FPGA-platform developed in the project. After exploring the design space with simulation, we will build a prototype in an FPGA.

**On-chip Network** Our previous research on design of on-chip networks targeted chip-implementations, and resulted in novel results on guaranteeing bandwidth in a clock-less setting [1, 2]. The circuit implementation is clever, but rather complex, and later work, including [5, 32], suggests that it is most likely a better option to aim for simpler and faster circuits. The higher speed will compensate for a poorer utilization of bandwidth. Concerning time predictability, our experience suggests that some form of time-slicing is preferable over solutions that try to make end-to-end guarantees in packet switched networks. This is believed to further simplify the circuit implementation. Our aim is to develop a simple, efficient and most likely (virtual) circuit switched on-chip network, which targets FPGA implementations.

The work will start with a study of possible alternative network architectures, with a focus on real-time issues, and a study of timing, wiring, and power consumption issues in big FPGA chips. This will be followed by design and implementation of an on-chip network and integration of this on-chip network into a complete multiprocessor platform. This research is well suited for a PhD project, and includes many opportunities for co-operation with other researchers and research groups. We have very good contacts to many of the key researchers in the field, and we have been selected to host the 6th ACM/IEEE Symposium on Networks-on-Chip 2012 at DTU.

**Processor Core** Time-predictable architectural features have been explored in the context of the Java processor JOP [24]. Caches are designed to be time-predictable and analyzable [23, 25]. Within this project we will adapt the time-predictable processor Patmos to FPGA technology. Patmos is currently under development in the EC project T-CREST and first concepts have already been published [29]. The VLIW design style of Patmos relies on static instruction schedules generated by the compiler in order to optimally exploit the available hardware resources. To support a time-predictable memory hierarchy (caches and scratchpad memories), Patmos will implement typed load and store instructions and prefetch instructions. We will adapt the LLVM compiler framework [17] to support the Patmos instruction set. Furthermore, we will integrate an open-source WCET analysis tool with the compiler to support WCET driven optimization.

**Prototype and Evaluation** The multiprocessor platform will be **implemented in FPGA technology** and one or more of the use-cases provided by Danfoss will be ported to this platform for a quantitative performance analysis, including **quantitative comparison with the existing design**, and verification.

**International Collaboration** We are the technical lead of the FP7 project T-CTRST which has many of the key players in Europe as partners and we expect a considerable synergy between the T-CREST project and the project proposed in this application. We also cooperate on research towards time-predictable architectures with Edward Lee's PRET group at the University of California, Berkeley (see e.g., [11]). One of our researchers will visit the PRET group at UCB for three months and we will host one PhD student from UCB for three months in summer 2013.

## 4 Dissemination and Publication Schedule

Scientific results will be published and presented at international conferences (DATE, NOCS, FPL, RTSS, DAC, CASES, CODES+ISSS) and in relevant scientific journals. One PhD thesis will publish the results from the project. We expect that each work package will result in several publications.

In the middle of the project and at the end of the project we will organize a workshop "FPGAs in Embedded Systems". Results from the project as well as projects in industry will be presented there. Given the growing use of FPGA technology in a wide and diverse range of Danish companies, such workshops represent timely attempts at bringing stakeholder's together and advancing the field.

To the extent possible, the results from the project will be available as open-source under the industry friendly BSD license. Open-source research projects attract other researchers, developers, and industrial partners to use and build on the results of the project. A project web site will host the project documentation, the published papers, and the source of the design.

## References


[1] T. Bjerregaard and J. Sparsø. A Router Architecture for Connection-Oriented Service Guarantees in the MANGO Clockless Network-on-Chip. In *Proc. Design Automation and Test in Europe (DATE)*, pages 1226–1231. IEEE Computer Society Press, 2005.

[2] T. Bjerregaard and J. Sparsø. A Scheduling Discipline for Latency and Bandwidth Guarantees in Asynchronous Network-on-chip. In *Proc. International Symposium on Asynchronous Circuits and Systems*, pages 34–43. IEEE Computer Society Press, 2005. (Best paper award).

[3] T. Bjerregaard and J. Sparsø. Implementation of Guaranteed Services in the MANGO Clockless Network-on-Chip. *IEE Proceedings: Computing and Digital Techniques*, 153(4):217–229, 2006.

[4] T. Bjerregaard and J. Sparsø. Packetizing OCP Transactions in the MANGO Network-on-Chip. In *9th EUROMICRO Conference on Digital System Design (DSD'06)*, pages pp. 657–664. IEEE Computer Society Press, 2006.

[5] T. Bjerregaard, M.B. Stensgaard, and J. Sparsø. A scalable, timing-safe, network-on-chip architecture with an integrated clock distribution method. In *Proc. Design Automation and Test in Europe (DATE)*, pages 648–653. IEEE Computer Society Press, 2007.

[6] Tobias Bjerregaard and Shankar Mahadevan. A survey of research and practices of network-on-chip. *ACM Computing Surveys*, 38(1):71–121, 2006.

[7] F. Clermidy, C. Bernard, R. Lemaire, J. Martin, I. Miro-Panades, Y. Thonnart, P. Vivet, and N. Wehn. A 477mw noc-based digital baseband for mimo 4g sdr. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, pages 278 – 279, February 2010.

[8] Christoph Cullmann, Christian Ferdinand, Gernot Gebhard, Daniel Grund, Claire Maiza, Jan Reineke, Benoît Triquet, and Reinhard Wilhelm. Predictability considerations in the design of multi-core embedded systems. In *Proceedings of Embedded Real Time Software and Systems*, May 2010.

[9] William Dally. Route packets, not wires: On-Chip interconnection networks. In *Proc. ACM/IEEE Design Automation Conference (DAC))*, pages 684–689, New York, June 2001. ACM Press.

[10] DaNES - Danish Network for Intelligent Embedded Systems. High-Technology Foundation, URL, http://www.danes.aau.dk/.

[11] Stephen A. Edwards, Sungjun Kim, Edward A. Lee, Isaac Liu, Hiren D. Patel, and Martin Schoeberl. A disruptive computer design idea: Architectures with repeatable timing. In *Proceedings of IEEE International Conference on Computer Design (ICCD 2009)*, Lake Tahoe, CA, October 2009. IEEE.

[12] Stephen A. Edwards and Edward A. Lee. The case for the precision timed (PRET) machine. In *DAC '07: Proceedings of the 44th annual conference on Design automation*, pages 264–265, New York, NY, USA, 2007. ACM.

[13] K. Goossens and A. Hansson. The Aethereal network on chip after ten years: Goals, evolution, lessons, and future. In *47th ACM/IEEE Design Automation Conference(DAC)*, pages 306 –311, June 2010.

[14] Trevor Harmon, Martin Schoeberl, Raimund Kirner, and Raymond Klefstad. A modular worst-case execution time analysis tool for Java processors. In *Proceedings of the 14th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2008)*, pages 47–57, St. Louis, MO, United States, April 2008. IEEE Computer Society.

[15] Reinhold Heckmann, Marc Langenbach, Stephan Thesing, and Reinhard Wilhelm. The influence of processor architecture on the design and results of WCET tools. *Proceedings of the IEEE*, 91(7):1038–1054, Jul. 2003.

[16] Thomas Henties, James J. Hunt, Doug Locke, Kelvin Nilsen, Martin Schoeberl, and Jan Vitek. Java for safety-critical applications. In *2nd International Workshop on the Certification of Safety-Critical Software Controlled Systems (SafeCert 2009)*, York, United Kingdom, Mar. 2009.

[17] Chris Lattner and Vikram S. Adve. LLVM: A compilation framework for lifelong program analysis & transformation. In *International Symposium on Code Generation and Optimization (CGO'04)*, pages 75–88. IEEE Computer Society, 2004.

[18] Isaac Liu, Jan Reineke, and Edward A. Lee. A PRET architecture supporting concurrent programs with composable timing properties. In *Signals, Systems and Computers, 2010 Conference Record of the Forty-Four Asilomar Conference on*, November 2010.

[19] Shankar Mahadevan, Federico Angiolini, Jens Sparsø, Luca Benini, and Jan Madsen. A Reactive and Cycle-True IP Emulator for MPSoC Exploration. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(1):109–122, January 2008.

[20] Christof Pitter and Martin Schoeberl. A real-time Java chip-multiprocessor. *ACM Trans. Embed. Comput. Syst.*, 10(1):9:1–34, 2010.

[21] M. S. Rasmussen, S. Karlsson, and J. Sparsø. Adaptable Support for Programming Models in Many-core Architectures. In *Workshop on New Directions in Computer Architecture*, 2009.

[22] RECOMP - Reduced Certification Costs Using Trusted Multi-core Platforms. ARTEMIS project. URL, http://atc.ugr.es/recomp/.

[23] Martin Schoeberl. A time predictable instruction cache for a Java processor. In *On the Move to Meaningful Internet Systems 2004: Workshop on Java Technologies for Real-Time and Embedded Systems (JTRES 2004)*, volume 3292 of *LNCS*, pages 371–382, Agia Napa, Cyprus, October 2004. Springer.

[24] Martin Schoeberl. A Java processor architecture for embedded real-time systems. *Journal of Systems Architecture*, 54/1–2:265–286, 2008.

[25] Martin Schoeberl. Time-predictable cache organization. In *Proceedings of the First International Workshop on Software Technologies for Future Dependable Distributed Systems (STFSSD 2009)*, pages 11–16, Tokyo, Japan, March 2009. IEEE Computer Society.

[26] Martin Schoeberl. Time-predictable computer architecture. *EURASIP Journal on Embedded Systems*, vol. 2009, Article ID 758480:17 pages, 2009.

[27] Martin Schoeberl, Florian Brandner, Jens Sparsø, and Evangelia Kasapaki. A statically scheduled time-division-multiplexed network-on-chip for real-time systems. In *Proceedings of the 6th International Symposium on Networks-on-Chip (NOCS)*, Lyngby, Denmark, May 2012. IEEE.

[28] Martin Schoeberl, Wolfgang Puffitsch, Rasmus Ulslev Pedersen, and Benedikt Huber. Worst-case execution time analysis for a Java processor. *Software: Practice and Experience*, 40/6:507–542, 2010.

[29] Martin Schoeberl, Pascal Schleuniger, Wolfgang Puffitsch, Florian Brandner, Christian W. Probst, Sven Karlsson, and Tommy Thorn. Towards a time-predictable dual-issue microprocessor: The Patmos approach. In *First Workshop on Bringing Theory to Practice: Predictability and Performance in Embedded Systems (PPES 2011)*, pages 11–20, Grenoble, France, March 2011.

[30] F. Steenhof, H. Duque, B. Nilsson, K. Goossens, and R.P. Llopis. Networks on chips for high-end consumer-electronics TV system architectures. In *Design, Automation and Test in Europe, 2006. DATE '06. Proceedings*, volume 2, pages 1–6, March 2006.

[31] M. B. Stensgaard, T. Bjerregaard, J. Sparsø, and J. H. Pedersen. A simple clockless Network-on-Chip for a commercial audio DSP chip. In *9th EUROMICRO Conference on Digital System Design (DSD'06)*, pages pp. 641–648. IEEE Computer Society Press, 2006.

[32] M. B. Stensgaard and J. Sparsø. ReNoC: A Network-on-Chip Architecture with Reconfigurable Topology. In *Proc. 2nd IEEE International Symposium on Networks-on-Chip*, pages 55–64. IEEE Computer Society Press, 2008.

[33] M. B. Stuart, M. B. Stensgaard, and J. Sparsø. Synthesis of Topology Configurations and Deadlock Free Routing Algorithms for ReNoC-based Systems-on-Chip. In *Proc. Embedded Systems Week 2009 - 7th IEEE/ACM International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, pages 481–490, 2009.

[34] T-CREST: Time-predictable Multi-Core Architecture for Embedded Systems. EU FP7-ICT Small or medium-scale focused research project (STREP). Duration: September 2011 to August 2014, (Partners: The Open Group, DTU Informatics, U. of York, T.U. Eindhoven, T.U. Vienna, AbsInt GMBH, GMV, Intecs S.p.a.). URL, http://cordis.europa.eu/projects/99768_en.html and http://www.3sei.com/t-crest/ .

[35] T. Ungerer, F. Cazorla, P. Sainrat, G. Bernat, Z. Petrov, C. Rochange, E. Quiñones, M. Gerdes, M. Paolieri, and J. Wolf. Merasa: Multi-core execution of hard real-time applications supporting analysability. *Micro, IEEE*, 30(5):66–75, 2010.

[36] Daniel Wiklund and Dake Liu. SoCBUS: Switched network on chip for hard real time embedded systems. *Parallel and Distributed Processing Symposium, International*, 0:78a, 2003.

[37] Reinhard Wilhelm, Daniel Grund, Jan Reineke, Marc Schlickling, Markus Pister, and Christian Ferdinand. Memory hierarchies, pipelines, and buses for future architectures in time-critical embedded systems. *IEEE Transactions on CAD of Integrated Circuits and Systems*, 28(7):966–978, 2009.