Contents lists available at ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

The ACROSS MPSoC – A new generation of multi-core processors designed for safety–critical embedded systems

Christian El Salloum *, Martin Elshuber, Oliver Höftberger, Haris Isakovic, Armin Wasicek

Vienna University of Technology, Institute of Computer Engineering, Treitlstrasse 3/3rd floor, 1040 Vienna, Austria

ARTICLE INFO

ABSTRACT

The European ARTEMIS ACROSS project aims to overcome the limitations of existing Multi-Processor Systems-on-a-Chip (MPSoC) architectures with respect to safety–critical applications. MPSoCs have a tremendous potential in the domain of embedded systems considering their enormous computational capacity and energy efficiency. However, the currently existing MPSoC architectures have significant limitations with respect to safety–critical applications. These limitations include difficulties in the certification process due to the high complexity of MPSoCs, the lacking temporal determinism and problems related to error propagation between subsystems. These limitations become even more severe, when subsystems of different criticality levels have to be integrated on the same computational platform. Examples of such mixed-criticality integration are found in the avionics and automotive industry with their desire to integrate safety–critical, mission critical and non-critical subsystems on the same platform in order to minimize size, weight, power and cost. The main objective of ACROSS is to develop a new generation of multi-core processors designed specially for safety–critical embedded systems; the ACROSS MPSoC. In this paper we will show how the ACROSS MPSoC overcomes the limitations of existing MPSoC architectures in order to make the multi-core technology available to the safety–critical domain. The proposed approach enables efficient certification, complexity management, mixed-criticality integration and the development of temporally deterministic hard real-time systems. The major technological innovations of ACROSS are an increased level of design abstraction, message-based interfaces for core-to-core communication and reliable fault and error containment established by a novel time-triggered network-on-chip. The achieved results comprise, a novel architecture for MPSoCs, a prototype implementation on FPGA technology as a proof-of-concept, a comprehensive set of middle-ware services and multiple demonstrators that show the benefits of the ACROSS Architecture in real world industrial applications.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

A Multi-Processor System-on-a-Chip (MPSoC) incorporates multiple, potentially heterogeneous processing cores and other functional units in a single case on a single die. Compared to general purpose single core processors, MPSoCs can provide enormous computational capacity in an energy efficient and cost efficient way. The roadmaps of the semiconductor industry [1] show a very clear trend towards multi-core technology, and we can safely assume that the majority of future high-end processors will be MPSoCs. Today, MPSoCs are typically applied in personal computers or consumer electronic devices like smart phones or tablets.

The scope of this paper is to elaborate MPSoCs in the domain of safety–critical embedded systems. Safety–critical systems are systems whose failure could result in loss of life, significant property damage, or damage to the environment. Examples are flight control systems for aircrafts, automotive control systems, medical devices, industrial control systems or nuclear power plants. As we will point out, MPSoCs could bring many benefits to safety–critical applications, but unfortunately, the currently existing MPSoC architectures were not designed with a strong focus on safety and certification and thus have serious drawbacks and limitations with respect to this domain. To overcome these limitations, a European consortium of 16 partners of industry and academia joined together in the ARTEMIS ACROSS project. A major result of this project is a new generation of multi-core processors designed specially for safety–critical embedded systems; the ACROSS MPSoC. In this paper we will show how the ACROSS MPSoC overcomes the limitations of existing MPSoC architectures in order to

* Corresponding author.

E-mail addresses: christian.el-salloum@tuwien.ac.at (C. El Salloum), martin.elshuber@tuwien.ac.at (M. Elshuber), oliver.hoeftberger@tuwien.ac.at (O. Höftberger), haris.isakovic@tuwien.ac.at (H. Isakovic), armin.wasicek@tuwien.ac.at (A. Wasicek).

make the multi-core technology available to the safety–critical domain.

The proposed approach enables efficient certification, complexity management, mixed-criticality integration and the development of temporally deterministic hard real-time systems. The major technological innovations of ACROSS are an increased level of design abstraction, message-based interfaces for core-to-core communication and reliable fault and error containment established by a novel time-triggered network-on-chip. The achieved results comprise, a novel architecture for MPSoCs, a prototype implemenation on FPGA technology as a proof-of-concept, a comprehensive set of middle-ware services and multiple demonstrators that show the benefits of the ACROSS Architecture in real world industrial applications.

The remainder of the paper is structured in the following way: Section 2 motivates the use of MPSoCs in embedded systems, while Section 3 points out problems and limitations of existing MPSoC architectures with respect to safety–critical systems. Section 4 outlines the related work. Section 5 introduces the ACROSS MPSoC architecture, of which we claim to overcome these limitations, and Section 6 describes a prototype implementation of the architecture. Section 7 presents an experimental validation of the MPSoC and Section 8 concludes the paper.

2. Benefits of MPSoCs in embedded systems

In order to motivate our work we start with showing the most important expected benefits of MPSoCs in the context of safety–critical embedded systems:

Energy and area efficiency. For decades, an increasing number of transistors was used to push the performance of a *single* processing core by developing larger micro-architectures with a higher complexity. Examples are micro-architectures based on super-pipelined designs featuring speculative, super-scalar, and out-of-order execution. The problem is that performance increases that are exclusively based on advances in micro-architecture are governed by *Pollack's Rule* [2]. Pollack's Rule states that, in the same process technology, a leading micro-processor consumes twice the area and power over the previous generation microprocessor, compared with a performance increase by a factor of 1.4. In other words, doubling the number of transistors in a single processor core results only in 40% additional performance, which means that ever increasing single core sizes yield diminishing performance in a power and area envelope. In contrast to the single core approach, a multi-core architecture has the potential to provide near linear performance improvement.

Computational performance. We have to be aware, that the potential speedup gained by a multi-core is limited by *Amdahl's Law*, which states that the theoretical maximum speedup achievable by parallelization is limited by the relative size of the non-parallelizable part (i.e. the serial part) of a program. Eq. (1) shows how the maximum speedup (S) relates to the relative size of the parallelizable part (p), the relative size of the serial part ($s = 1 - p$), and the number of parallel cores (n).

$$S = \frac{s + p}{s + p/n} = \frac{1}{s + (1-s)/n} \qquad (1)$$

The formula shows that even a small percentage of non-parallelizable code in a program leads to a saturation of the achievable speedup at a small number of cores.

It must be considered that this limitation is only valid if one tries to parallelize a single application with a single continuous serial part across all the cores in the chip. In the domain of embedded systems, a device typically integrates multiple application subsystems that are inherently parallel. As an example consider the electronic control system in a modern vehicle executing tasks for the power train, the comfort electronics, or for the vehicle dynamics management system. In such a system the real challenge is not to parallelize algorithms solving *one* big problem as it is done for high-performance computing, but to efficiently integrate multiple tasks that are inherently parallel. So what we want to point out here, is that Amdahl's Law [3] which is considered as the most limiting factor concerning parallelization for high-performance computing is not necessarily a significant constraint in many embedded applications.

Heterogeneity. A typical embedded application consists of multiple subsystems and tasks with different requirements to the underlying hardware. Examples are encoding and decoding of signals like video or audio streams, encryption and decryption tasks in order to establish required security properties, or data transformations for control loops. Of course, all these tasks could be theoretically executed on general purpose CPUs, but in practice this is very inefficient and often infeasible due to the strict time and energy constraints of embedded applications. Good examples are high frequency control loops where specialized hardware like Digital Signal Processors (DSPs) is required to meet the specified deadlines, or security hardware extensions that significantly increase the energy efficiency of the implementation of a cipher.

Constructing a single core that strives to satisfy all possible kinds requirements, always result in a more or less inefficient trade-off solution. MPSoCs enable the combination of multiple heterogeneous cores on a single die that can be highly optimized for the individual tasks in an embedded application. This approach is already followed in many consumer devices, where dedicated cores with different functionalities (e.g., graphic processors, wireless communication components, general purpose CPUs) are integrated to form an efficient System-on-a-Chip.

Reduction of physical units. The high computational capacity and the possibility to integrate multiple heterogeneous cores on a single die, make an MPSoC optimally suitable for integrating multiple functionalities of an embedded system into a single chip, and thus into a single device. In many embedded applications, the consolidation of multiple functionalities in a single device can lead to massive savings with respect to cost, energy, volume and weight.

Modern cars already contain up to 70 Electronic Control Units (ECU), which are interconnected via multiple different communication networks. By integrating multiple functionalities in the ECUs, the total number of ECUs can be reduced, which directly leads to lower manufacturing costs since fewer cases, electronic parts, cables, connectors and assembly steps on the production line are required.

Beside these savings, a reduction of the numbers of ECUs can lead to improved reliability of the entire system. The rationale is that if there are fewer components in the system, there are also fewer components that can fail. In particular this is true regarding the involved reduction of connectors. By analyzing field data from the automotive industry it was shown that more than 30% of electrical failures are caused by connector problems [4].

3. Limitations of existing MPSoC architectures with respect to safety–critical applications

Since MPSoCs can bring many benefits to the domain of embedded systems, developers and manufacturers are eager to employ this technology for their products. However, existing MPSoC architectures have significant limitations with respect to safety–critical applications. In this section we point out these limitations and explain why current architectures are designed that way.