

Preface

The ever increasing complexity and integration densities of multiprocessor system-on-chip (MPSoC) architectures, significantly enlarges the design space of embedded computing systems. A wide range of design choices must be tuned from a multi-objective perspective, mainly in terms of performance and energy consumption, to find the most suitable system configuration for the target application. Given the huge design space to be analysed, the exploration of tradeoffs between multiple competing objectives cannot be anymore driven by a manual optimisation process based on the intuition and past experience of the system architect. Multi-objective exploration of the huge design space of next generation multi/many core architectures needs for Automatic Design Space Exploration techniques to systematically explore the design choices and to compare them in terms of multiple competing objectives (trade-offs analysis). Ideally, a designer would try all possible design choices and choose the most suitable according to the specific system requirements. Unfortunately, such an exhaustive approach is often unfeasible because of the large number of design choices to be simulated and analysed, in some cases showing some sophisticated effects on system properties that rarely enable to easily and accurately model the system behavior. Consequently, good search techniques are needed not only to find design alternatives that best meet system constraints and cost criteria, but also to prune the search space to crucial parameters to enable an effective and efficient design space exploration.

In the age of multi/many core architectures, system optimization and exploration definitely represent challenging research tasks. Although many point tools exist to optimize particular aspects of embedded systems, an overall design space exploration framework is needed to combine all the decisions into a global search space with a common interface to the optimization and evaluation tools. The state-of-the art lacks of a Design Space Exploration (DSE) framework to help the designer in the automatic selection of the most suitable system configuration for a certain application given a set of multiple competing objectives.

Based on the idea to provide an automatic DSE framework, we started thinking about a proposal of an European project, namely MULTICUBE¹, that was submitted in May 2007 in the first call of the Seventh Framework Programme on ICT under the Objective 3.3 on Embedded Systems Design. The proposal was accepted and the project started on January 2008 under the coordination of Politecnico di Milano. This book was mainly catalyzed by the main research outcomes and exploitable results of the MULTICUBE European project, where the Editors have acted in 2008–2010 timeframe. The MULTICUBE project focused on the definition of an automatic multi-objective Design Space Exploration (DSE) framework to be used to tune the System-on-Chip architecture for the target application evaluating a set of metrics (e.g. energy, latency, throughput, bandwidth, QoR, etc.) for the next generation embedded multimedia platforms.

The project aimed at increasing the competitiveness of European industries by optimizing the design of embedded computing systems while reducing design time and costs. The project defined an automatic multi-objective DSE framework to be used at design-time to find the best power/performance trade-offs while meeting system-level constraints and speeding up the exploration process. A set of heuristic optimization algorithms have been defined to reduce the exploration time, while a set of response surface modeling techniques have been defined to further speed up the process. Based on the results of the design-time multi-objective exploration, the MULTICUBE project also defined a methodology to be used at run-time to optimize the allocation and scheduling of different application tasks. The design exploration flow results in a Pareto-optimal set of design alternatives in terms of power/performance trade-offs. This set of operating points can then be used at run-time to decide how the system resources should be distributed over different application tasks running on the multiprocessor system on chip.

The MULTICUBE DSE framework leverages a set of open-source and proprietary tools for the exploration, modeling and simulation to guarantee a wide exploitation of the MULTICUBE project results in the embedded system design community. The integration of different tools is ensured by a common XML-based tool interface specification, defined to enable the independent development of modules and a seamless integration of the design tools and the data structures into a common design environment. Several industrial use cases (defined as combination of application and related architecture) have been used to assess the capabilities of the MULTICUBE design flow and tools in an industrial design process. The MULTICUBE project has been strongly industry-driven: industrial partners (STMicroelectronics and DS2) as well IMEC research center have defined the design techniques and tools requirements and then validated them to design some industrial use cases. The benefits of the introduction of the automatic DSE in the design phase of embedded computing systems (justifying its introduction in industrial design processes) have been assessed through a procedure to assess the final objective design quality and the reduction of design turnaround time by introducing such a technology on the entire design

¹ The project acronym, MULTICUBE, stands for: “Multi-Objective design space exploration of multi-processor SoC architectures for embedded multimedia applications”.

process. The benefits on the design process can be measurable and tangible like the reduction of the overall design process lead time, and qualitative or intangible like the streamlining and the reduction of human error prone repetitive operations. The DSE assessment procedure was the basis for the validation of the industrial use cases and demonstrators of the project. Validation results have been assessed based on a set of common assessment criteria.

In the book, we have tried to provide a comprehensive understanding of several facets of the problem of design space exploration for embedded on-chip architectures. The book chapters are organized in two parts. In Part I, several methodologies and tools to support automatic design space exploration are discussed. In Part II of the book, the DSE methodologies and tools described in Part I have then been applied to several application domains to discuss their applicability and to envision their benefits. First, a high-level modeling and exploration approach has been applied for a powerline communication network based on a SoC, then the application of the automatic DSE flow to parallel on-chip architectures is discussed, and finally the DSE for run-time management has been applied to a reconfigurable system for video streaming.

Entering Part I on methodologies and tools, Chap. 1 introduces the MULTICUBE design-flow to support the automatic multi-objective Design Space Exploration (DSE) to tune the parameters of System-on-Chip architectures by considering several metrics such as energy, latency and throughput. One of the important goals of the DSE framework is to find design trade-offs that best meet the system constraints and the cost criteria which are indeed strongly dependent on the target application. The DSE flow is based on the interaction of two frameworks to be used at design time: the Design Space Exploration Framework, a set of open-source and proprietary architectural exploration tools, and the Power/Performance Estimation Framework, a set of modeling and simulation tools (open-source and proprietary) operating at several levels of abstraction. The DSE flow also includes the specification of an XML integration interface to connect the exploration and estimation frameworks and a Run-time Resource Manager that selects, at run-time, the best software configuration alternatives to achieve a good power/performance trade-off.

Chapter 2 introduces M3-SCoPE, an open-source SystemC-based framework for performance modeling of multi-processor embedded systems, software source code behavioral simulation and performance estimation of multi-processor embedded systems. Using M3-SCoPE, the application software running on the different processors of the platform can be simulated efficiently in close interaction with the rest of the platform components. In this way, fast and accurate performance metrics of the system are obtained. These metrics are then delivered to the DSE tools to evaluate the quality of the different configurations in order to select the best power/performance trade-offs.

Chapter 3 presents the optimization algorithms developed in the MULTICUBE project for Design Space Exploration of embedded computing systems. Two software DSE tools implement the optimization algorithms: M3Explorer and modeFRONTIER. The mathematical details of the given optimization problems are explained in

the chapter together with how the algorithms can exchange information with the simulators. The description of the proposed algorithms is the central part of the chapter. The focus is posed on new algorithms and on “ad hoc” modifications implemented in existing techniques to face with discrete and categorical variables, which are the most relevant ones when dealing with embedded systems design. The strategy to test the performance achieved by the optimization is another important topic treated in the chapter. The aim is mainly to build confidence in optimization techniques, rather than to simply compare one algorithm with respect to another one. The “no-free-lunch theorem for optimization” has to be taken into consideration and therefore the analysis will look forward to robustness and industrial reliability of the results. The main contribution of MULTICUBE project in the research field of optimization techniques for embedded systems design is indeed the high level of the obtained compromise between specialization of the algorithms and concrete usability of the DSE tools.

A typical design space exploration flow involves an event-based simulator in the loop, often leading to an actual evaluation time that can exceed practical limits for realistic applications. Chip multi-processor architectures further exacerbate this problem given that the actual simulation speed decreases by increasing the number of cores of the chip. Traditional design space exploration lacks of efficient techniques that reduce the number of architectural alternatives to be analyzed. In Chap. 4, we introduce a set of Response Surface Modeling (RSM) techniques that can be used to predict system level metrics by using closed-form analytical expressions instead of lengthy simulations. The principle of RSM is to exploit a set of simulations generated by one or more Design of Experiments strategies to build a surrogate model to predict the system-level metrics. The response model has the same input and output features of the original simulation based model but offers significant speed-up by leveraging analytical, closed-form functions which are tuned during a model training phase.

Running multiple applications optimally in terms of Quality of Service (e.g., performance and power consumption) on embedded multi-core platforms is a huge challenge. Moreover, current applications exhibit unpredictable changes of the environment and workload conditions which makes the task of running them optimally even more difficult. Chapter 5 presents an automated tool flow which tackles this challenge by a two-step approach: first at design-time, a Design Space Exploration (DSE) tool is coupled with a platform simulator(s) to get optimum operating points for the set of target applications. Secondly, at run-time, a lightweight Run-time Resource Manager (RRM) leverages the design-time DSE results for deciding an operating configuration to be loaded at run-time for each application. This decision is taken dynamically, by considering the available platform resources and the QoS requirements of the specific use-case. To keep RRM execution and resource overhead at minimum, a very fast optimisation heuristic is integrated demonstrating a significant speedup in the optimisation process, while maintaining the desired Quality of Service.

Emerging MPSoC platforms provide the applications with an extended set of physical resources, as well as a well a defined set of power and performance optimization mechanisms (i.e., hardware control knobs). The software stack, meanwhile,

is responsible of taking directly advantage of these resources, in order to meet application functional and non-functional requirements. The support from the Operating System (OS) is of utmost importance, since it gives opportunity to optimize the system as a whole. The main purpose of Chap. 6 is to introduce the reader to the challenges of managing physical and logical resources in complex multi/many-core architectures at the OS level.

Entering Part II on application domains, Chap. 7 presents the application of MULTICUBE methodology to the design of an ITU G.hn compatible component for a powerline communication network based on a SoC. Powerline communication is an advanced telecommunication system enabling fast and reliable transfer of audio, video and data information using the most ubiquitous transmission system: the power lines. This transmission line is used to exchange information between different equipment connected to the network using the advanced coding techniques like such as Orthogonal Frequency Division Multiplexing. The starting point of the analysis is a high level SystemC-based virtual platform for which the chapter analyzes the effects of the variation of a pre-defined set of design parameters on a set of pre-defined metrics. This automatic analysis will drive the design choices in order to build an optimized industrial system. The chapter shows that the SystemC-based virtual platform combined with the MULTICUBE design space exploration framework can save up to 80% of designer's work time, while achieving better results in terms of performance.

Chapter 8 describes two significant applications of the automatic MULTICUBE DSE flow to parallel on-chip architectures. The first part of the chapter presents the design space exploration of a low power processor developed by STMicroelectronics by using the modeFRONTIER tool to demonstrate the benefits DSE not only in terms of objective quality, but also in terms of impact on the design process within the corporate environment. The second part of the chapter describes the application of Response Surface Models introduced in Chap. 4 to a tiled, multiple-instruction, many-core architecture developed by the Chinese Academy of Sciences. Overall, the results have showed that different models can present a trade-off of accuracy versus computational effort. In fact, throughout the evaluation, we observed that high accuracy models require high computational time (for both model construction time and prediction time); vice-versa low model construction and prediction time has led to low accuracy.

Chapter 9 reports a case study of DSE for supporting Run-time Resource Management (RRM). The management of system resources for an MPSoC dedicated to multiple MPEG4 encoding is addressed in the context of an Automotive Cognitive Safety System. The runtime management problem is defined as the minimization of the platform power consumption under resource and Quality of Service (QoS) constraints. The chapter provides an insight of both, design-time and run-time aspects of the problem. During the preliminary design-time DSE phase, the best configurations of run-time tunable parameters are statically identified for providing the best trade-offs in terms of run-time costs and application QoS. To speed up the optimization process without reducing the quality of final results, a multi-simulator framework is used for modeling platform performance. At run-time, the RRM exploits the

design-time DSE results for deciding an operating configuration to be loaded for each MPEG4 encoder. This operation is carried out dynamically, by following the QoS requirements of the specific use-case.

Due to the large number of topics discussed in the book and their heterogeneity, the background on system modeling, simulation and exploration is discussed chapter by chapter with a separate reference set for each chapter. This choice also contributed to make each chapter self-contained.

Overall, we believe that the book chapters cover a set of definitely important and timely issues impacting the present and future research on automatic DSE for embedded multi-core on-chip architectures. We sincerely hope that the book could become a solid reference in the next years. In our vision, the authors put a big effort in clearly presenting their technical contributions outlining the potential impact and benefits of the proposed approach on same case studies. Our warmest gratitude goes to the MULTICUBE team, for their continuous effort and dedication during the project and for their contribution as authors of the chapters. We would like to gratefully acknowledge our EC Project Officer, Panagiotis Tsarchopoulos and our EC Project Reviewers: Alain Perbost, Andrzej Pulka and Kamiar Sehat for their valuable comments and guidance during the MULTICUBE project. A special thanks to Charles Glaser from Springer for encouraging us to write a single textbook on the topic of design space exploration based on our experience of the MULTICUBE project.

With our work on MULTICUBE project and this book, we have pushed towards the adoption of automatic design space exploration for the design of multi-processor architectures for embedded computing systems. This book is expected to be one of the most important dissemination vehicles to spread out the knowledge developed in the MULTICUBE project in the international community after the end of the project.

Milano, Italy
Milano, Italy
Santander, Spain
March 2011

The Editors,
Cristina Silvano
William Fornaciari
Eugenio Villar



<http://www.springer.com/978-1-4419-8836-2>

Multi-objective Design Space Exploration of Multiprocessor
SoC Architectures

The MULTICUBE Approach

(Eds.) C. Silvano; W. Fornaciari; E. Villar

2011, XVIII, 222 p. 88 illus., Hardcover

ISBN: 978-1-4419-8836-2