

UNIVERSITY OF SOUTHAMPTON

**ENERGY MINIMISATION TECHNIQUES
FOR DISTRIBUTED EMBEDDED SYSTEMS**

by

Marcus Thomas Schmitz

A thesis submitted for the degree of
Doctor of Philosophy

Department of Electronics and Computer Science,
Faculty of Engineering and Applied Science,
University of Southampton,
United Kingdom.

February 2003

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND APPLIED SCIENCE
DEPARTMENT OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

Energy Minimisation Techniques for Distributed Embedded Systems

by Marcus Thomas Schmitz

It is likely that the demand for embedded computing systems with low energy dissipation will continue to increase. This thesis is concerned with the development and validation of techniques that allow an effective automated design of energy-efficient embedded systems. Special emphasis is placed upon system-level co-synthesis techniques for systems that contain dynamic voltage scalable processors which can trade off between performance and power consumption during run-time.

The first part of the thesis addresses energy minimisation of distributed embedded systems through dynamic voltage scaling (DVS). A new voltage selection technique for single-mode systems based on a novel energy-gradient scaling strategy is presented. This technique, which overcomes limitations of previous voltage scaling approaches, exploits system idle and slack time to reduce the power consumption, taking into account the individual task power dissipation. Numerous benchmark experiments validate the quality of the proposed technique in terms of energy reduction and computational complexity.

The second part of the thesis focuses on the development of genetic algorithm-based co-synthesis techniques (mapping and scheduling) for single-mode systems that have been specifically developed for an effective utilisation of the voltage scaling approach introduced in the first part. The schedule optimisation improves the execution order of system activities not only towards performance, but also towards a high exploitation of voltage scaling to achieve energy savings. The mapping optimisation targets the distribution of system activities across the system components to further improve the utilisation of DVS, while satisfying hardware area constraints. Extensive experiments including a real-life optical flow detection algorithm are conducted, and it is shown that the proposed co-synthesis techniques can lead to high energy savings with moderate computational overhead.

The third part of this thesis concentrates on energy minimisation of emerging distributed embedded systems that accommodate several different applications within a single device, i.e., multi-mode embedded systems. A new co-synthesis technique for multi-mode embedded systems based on a novel operational-mode-state-machine specification is presented. The technique increases significantly the energy savings by considering the mode execution probabilities that yields better resource sharing opportunities.

The new co-synthesis and voltage scaling techniques have been incorporated into the prototype co-synthesis tool LOPOCOS (Low Power Co-Synthesis). The capability of LOPOCOS in efficiently exploring the architectural design space is demonstrated through a system-level design of a realistic smart phone example that integrates a GSM cellular phone transcoder, an MP3 decoder, as well as a JPEG image encoder and decoder.

Acknowledgements

First of all I would like to express my gratitude to my supervisor Dr. Bashir M. Al-Hashimi, whose constant support and encouragement throughout the last three (and a half) years have been of great benefit. He taught me how to clearly structure my ideas and how to write them down into readable technical papers.

I would like to thank the members of the Electronic Systems Design Group (ESD) at the University of Southampton. Especially, I wish to express sincere thanks to Professor Andrew D. Brown for constructive feedback during my MPhil to PhD transfer examination. I am also thankful to my colleagues Theo Gonciari, Neil Henderson, Nicola Nicolici, Paul Rosinger, and Mauricio Varea for their omnipresent friendship and for establishing an inspiring cross-cultural working atmosphere.

I am deeply grateful to Professor Petru Eles from Linköping University (Sweden) for many valuable successions and excellent comments, as well as for providing financial support during an one-month stay at Linköping University.

Financial support of my work at the University of Southampton was provided by the Department of Electronics and Computer Science as well as by the Engineering and Physical Sciences Research Council (EPSRC). I would like to acknowledge Neal K. Bambha (University of Maryland, USA) and Flavius Gruian (Lund University, Sweden) for kindly providing their benchmark sets.

Last, but most importantly, I have to and I wish to acknowledge the people which are the most closest to me. My parents Doris and Werner, my fiancée Elsa, and my brother Christian have been the stabilising forces in my life. I dedicate this thesis to them. It will be hard to return all the unlimited support and unconditional love which I have received. Elsa's liveliness and love have kept my mind away from any drowsiness, especially when lying on the ground observing shooting stars in the middle of the night. She embodies a daily source of happiness and bliss. Christian has allowed helped in deriving the smart phone task graphs during a short stay in Southampton.

Table of Contents

Abstract	i
List of Figures	ix
List of Tables	xi
Abbreviations	xiii
1 Introduction	1
1.1 Embedded System Design Flow	2
1.2 System Specification	5
1.2.1 Task Graph Representation	6
1.3 Co-Synthesis	8
1.3.1 Architecture Allocation	9
1.3.2 Application Mapping	10
1.3.3 Activity Scheduling	12
1.3.4 Energy Management	13
1.4 Hardware and Software Synthesis	15
1.5 Contributions and Thesis Overview	19
2 Background and Previous Work	22
2.1 Energy Dissipation of Processing Elements	23
2.2 Energy Minimisation Techniques	28
2.2.1 Dynamic Voltage Scaling	28
2.2.2 Dynamic Power Management	29
2.2.3 Dynamic Power Management versus Dynamic Voltage Scaling . .	30
2.2.4 DVS and DPM in Distributed Embedded Systems	31

2.3	Energy Dissipation of Communication Links	33
2.4	Previous Work	34
2.4.1	Co-Synthesis without Energy Minimisation	34
2.4.2	Co-Synthesis with Energy Minimisation	35
2.5	Concluding Remarks	36
3	Power Variation-Driven Voltage Scaling	37
3.1	Motivation	38
3.2	Novel Algorithms for Dynamic Voltage Scaling	46
3.2.1	Energy-Gradient-Based Voltage Scaling	47
3.2.2	Discrete Voltages	51
3.2.3	Algorithm Complexity	51
3.3	Experimental Results: Energy-Gradient based Dynamic Voltage Scaling .	52
3.3.1	Power Variations	54
3.3.2	Minimal Extension Time	58
3.4	Concluding Remarks	62
4	Optimisation of Mapping and Scheduling for Dynamic Voltage Scaling	63
4.1	Schedule Optimisation	64
4.1.1	Motivational Example: Scheduling	65
4.1.2	Background	68
4.1.3	Genetic List Scheduling Algorithm	70
4.1.4	Experimental Results: Schedule Optimisation	79
4.2	Optimisation of Task and Communication Mapping	84
4.2.1	Genetic Task Mapping Algorithm	85
4.2.2	Combined Scheduling and Communication Mapping	88
4.2.3	Experimental Results: Mapping Optimisation	92
4.3	Optimisation of Allocation	98
4.3.1	Experimental Results: Component Allocation	99
4.4	Concluding Remarks	101
5	Energy-Efficient Multi-Mode Embedded Systems	103
5.1	Preliminaries	104
5.1.1	Functional Specification of Multi-Mode Systems	104
5.1.2	Architectural Model and System Implementation	107

5.2	Motivational Examples	109
5.3	Previous Work	112
5.4	Problem Formulation	114
5.5	Co-Synthesis of Energy-Efficient Multi-Mode Systems	116
5.5.1	Multi-Mode Co-Synthesis Algorithm	116
5.5.2	Hardware Core Allocation	122
5.5.3	Dynamic Voltage Scaling for Multiple Parallel Executing Tasks	122
5.6	Experimental Results: Multi-Mode	126
5.6.1	Hypothetical Examples	126
5.6.2	Smart Phone Benchmark	132
5.7	Concluding Remarks	134
6	LOPOCOS: A Prototype Low Power Co-Synthesis Tool	136
6.1	Smart Phone Description	137
6.1.1	Voice Compression and Decompression	137
6.1.2	MP3 Decoder	141
6.1.3	JPEG Image Compression and Decompression	141
6.2	LOPOCOS	143
6.2.1	Input Descriptions	145
6.2.2	Architectural Design Space Exploration	151
6.3	Concluding Remarks	158
7	Conclusion	160
7.1	Summary and Research Contributions	161
7.2	Future Research Directions	164
7.2.1	Control-flow Intensive Applications	164
7.2.2	Energy-aware Granularity Selection	164
7.2.3	Platform Identification	165
A	Supplementary Experiments	166
B	Smart Phone Benchmark	170
B.1	GSM Voice Compression	170
B.2	GSM Voice Decompression	173
B.3	MP3 Sound Decoder	174

B.4	JPEG Image Encoder	175
B.5	JPEG Image Decoder	176
B.6	Technology File	176
C	DVS Problem Formulation for Distributed Heterogeneous Systems	184
D	Derivation of the Energy/Delay Trade-Off	187
	References	188

List of Figures

1.1	Example of a typical embedded system (smart-phone)	2
1.2	Typical design flow of a new embedded computing system	4
1.3	MP3 decoder given as (a) task graph specification (17 tasks and 18 communications) and (b) high-level language description in C	7
1.4	System-level co-synthesis flow	8
1.5	Architectural selection problem	9
1.6	Application mapping onto hardware and software components	11
1.7	Two different scheduling variants based on the same allocated architecture and identical application mapping	13
1.8	System schedule with idle and slack times	14
1.9	The concept of dynamic voltage scaling	15
1.10	Hardware synthesis flow	17
1.11	Software synthesis flow	18
2.1	Dynamic power dissipation of an inverter circuit [34]	24
2.2	Supply voltage dependent circuit delay	26
2.3	Energy versus delay function using fixed and dynamic supply voltages (considering $V_{max} = 3.3V$ and $V_t = 0.8V$)	27
2.4	Block diagram of DVS-enabled processor [33]	28
2.5	Shutdown during idle times (DPM)	31
2.6	Voltage scaling to exploit the slack time (DVS)	31
2.7	Combination of dynamic voltage scaling and dynamic power management	32
3.1	Architecture and specification for the motivational example	39
3.2	Power profile of a possible mapping and schedule at nominal supply voltage (no DVS is applied)	41
3.3	Two different voltage scaled schedules	44

3.4	Pseudo code of the proposed heuristic (PV-DVS) algorithm	48
3.5	Capturing the mapping and schedule information into the task graph by using pseudo edges and communication task	49
3.6	Pseudo code of task graph to mapped-and-scheduled task graph transfor- mation	49
3.7	Three identical execution orders of the <code>tgff17_m</code> benchmark: (a) un- scaled execution at nominal supply voltage (NO-DVS), (b) using the EVEN- DVS, and (c) the PV-DVS approach	57
3.8	Energy reduction quality dependent on minimal extension time Δt_{min} . . .	58
3.9	Execution time dependent on minimal extension time Δt_{min}	59
3.10	Energy reduction quality dependent on execution time	60
4.1	Co-synthesis flow for the optimisation of scheduling and mapping to- wards the utilisation of PV-DVS	64
4.2	Specification and DVS-enabled architecture	65
4.3	A possible schedule <i>not</i> optimised for DVS	66
4.4	Schedule optimised for DVS considering the power variation model . . .	67
4.5	List scheduling	72
4.6	Task priority encoding into a priority string	73
4.7	Principle behind the genetic list scheduling algorithm	74
4.8	Proposed EE-GLSA approach for energy-efficient schedules	76
4.9	Hole filling problem	77
4.10	Task mapping string describing the mapping of five tasks to an architecture	86
4.11	Proposed EE-GTMA approach for energy-efficient task mappings	87
4.12	Combined optimisation of task and communication mapping	88
4.13	A combined priority and communication mapping string	90
4.14	Proposed EE-GLSCMA approach for combined optimisation of energy- efficient schedules and communication mappings	91
4.15	Three scheduling and mapping concepts	93
4.16	Nine different implementation possibilities of the OFD algorithm	101
5.1	Example operational mode state machine of a smart phone	105
5.2	Relation between OMSM and individual task graph specifications	106
5.3	Distributed Architectural Model	108
5.4	Mode execution probabilities	109

5.5	Multiple task type implementations	112
5.6	Task mapping string for multi-mode systems	117
5.7	Pseudo Code: Multi-Mode Co-Synthesis	118
5.8	Pseudo Code: Mapping Modification towards component shutdown	121
5.9	DVS Transformation for HW Cores	123
5.10	DVS Transformation for HW Cores considering inter-PE communication	124
5.11	Pseudo code: Task graph transformation for DVS-enabled hardware cores	125
5.12	Pareto optimal solution space achieved through a single optimisation run of mul115 (without DVS), revealing the solution trade-offs between energy dissipation and area usage	130
5.13	A system specification consisting of two operational modes optimised for different execution probabilities (red=0.1:0.9, blue=0.9:0.1, green=0.5:0.5)	131
5.14	Energy dissipation of the Smart phone using different optimisation strategies	134
6.1	Block diagram of the GSM RPE-LTP transcoder [69]	138
6.2	Task graph of the GSM voice encoder	139
6.3	Task graph of the GSM voice decoder	140
6.4	Block diagram of the MPEG-1 layer 3 audio decoder	141
6.5	Block diagram of the JPEG encoder and decoder [142]	142
6.6	Task graphs of the JPEG encoder and decoder	142
6.7	Design flow used within LOPOCOS	144
6.8	File description of the top-level finite state of the smart phone	146
6.9	File description of a single mode task graph	147
6.10	Technology library file	149
6.11	Co-synthesis results of Architecture 1	154
6.12	Co-synthesis results of Architectures 2 and 3	155
6.13	Co-synthesis results for Architectures 2 and 3, exploiting DVS	157
6.14	Co-synthesis results for Architecture 4	158

List of Tables

1.1	Trade-offs between serveral heterogeneous components (+ + \equiv highly advantageous, + \equiv advantageous, o \equiv moderate, - \equiv disadvantageous, - - \equiv highly disadvantageous)	10
1.2	Task execution properties (time and power) on different processing elements	11
3.1	Nominal task execution times and power dissipations	40
3.2	Communication times and power dissipations of communication activi- ties mapped to the bus	40
3.3	Evolution of the energy-gradients during voltage scaling	45
3.4	Comparison of the presented PV-DVS optimisation with the fixed power model using EVEN-DVS approach	55
3.5	PV-DVS results using the benchmarks of Bambha <i>et al.</i> [19]	56
4.1	Nominal execution times and power dissipations for the mapped tasks . .	66
4.2	Experimental results obtained using the <i>fixed power model</i> and the <i>power variation model</i> during voltage selection; both integrated into a genetic list scheduling algorithm	81
4.3	Experimental results obtained using the generalised DVS optimised schedul- ing approach for benchmark examples TG1 and TG2	83
4.4	Mapping optimisation with and without DVS optimised scheduling using tgff and hou benchmarks	94
4.5	Mapping optimisation of the benchmark set TG1 using NO-DVS (Nomi- nal), EVEN-DVS, and PV-DVS	96
4.6	Comparison between DLS algorithm and the proposed scheduling and mapping approach using Bambha's benchmarks [19]	97
4.7	Increasing architectural parallelism to allow voltage scaling of the OFD algorithm	100

4.8	Relaxing the performance constraints of the OFD algorithm	100
5.1	Task execution and implementation properties	110
5.2	Considering mode execution probabilities (excluding DVS)	127
5.3	Considering mode execution probabilities (including DVS)	129
5.4	Results of smart phone experiments	133
6.1	Task independent components parameters	150
6.2	Task dependent parameters	151
6.3	Components in a typical technology library	152
A.1	Profiling results of music piece I	167
A.2	Profiling results of music piece II	168
A.3	Profiling results of music piece III	169

Abbreviations

ASAP	As soon as possible
ALAP	As late as possible
ALU	Arithmetic logic unit
APM	Advanced power management
ASIC	Application specific integrated circuit
ASIP	Application specific instruction processor
CAD	Computer-aided design
CDFG	Control-data flow graph
CMOS	Complementary metal-oxide semiconductor
CORDIC	Coordinate rotation digital computing
CPU	Central processing unit
DAG	Directed acyclic graph
DSP	Digital signal processor
DHS	Distributed heterogeneous system
DVS	Dynamic voltage scaling
DVS-PE	Dynamic voltage scalable processing element
EDA	Electronic design automation
EPST	Earliest possible start time
FIFO	First in first out
FIR	Finite impulse response
FPGA	Field programmable gate array
FPU	Floating-point unit
FSM	Finite state machine
GA	Genetic algorithm
GLSA	Genetic list scheduling algorithm
GPP	General purpose processor

HDL	Hardware description language
HW	Hardware
IDCT	Inverse discrete cosine transformation
ILP	Integer linear programming
LCM	Least common multiplier
LS	List scheduling
MEMS	Mirco electro mechanical systems
MSTG	Mapped and scheduled task graph
OMSM	Operational mode state machine
PDA	Personal digital assistant
PE	Processing element
PV-DVS	Power variation dynamic voltage scaling
QoS	Quality of Service
RAM	Random access memory
ROM	Read-only memory
RTL	Register-transfer level
RTOS	Real-time operating system
RTS	Real-time system
SA	Simulated annealing
SoC	System on a chip
SRAM	Static RAM
SW	Software
TS	Tabu search
VCO	Voltage controlled oscillator
VLSI	Very large scale integration
WCET	Worst case execution time

Chapter 1

Introduction

Over the last several years, the popularity of portable applications has explosively increased. Millions of people use battery-powered mobile phones, digital cameras, MP3 players, and personal digital assistants (PDAs). To perform major parts of the system's functionality, these mass products rely, to a great extent, on sophisticated embedded computing systems with *high performance* and *low power dissipation*. The complexity of such devices, caused by an ever-increasing demand for functionality and feature richness, has made the design of modern embedded systems a time-consuming and error-prone task. To be commercially successful in a highly competitive market segment with tight time-to-market and cost constraints, computer-based systems in mobile applications should be cheap and quick to realise, while, at the same time, consume only a small amount of electrical power, in order to extend the battery-lifetime. Designing such embedded systems is a challenging task.


This thesis addresses this problem by providing techniques and algorithms for the automated design of energy-efficient distributed embedded systems which have the potential to overcome traditional design techniques that neglect important *energy management* issues. In this context, special attention is drawn to *dynamic voltage scaling* (DVS) — an energy management technique. The main idea behind DVS is to dynamically scale the supply voltage and operational frequency of digital circuits during run-time, in accordance to the temporal performance requirements of the application. Thereby, the energy dissipation of the circuit can be reduced by adjusting the system performance to an appropriate level. Furthermore, the proposed synthesis techniques target the coordinated design

(co-design) of *mixed hardware/software* applications towards the effective exploitation of DVS, in order to achieve substantial reductions in energy.

The main aims of this chapter are to introduce the fundamental problems that are involved in designing distributed embedded systems and to provide the terminology used throughout this work. The remainder of this chapter is organised as follows. Section 1.1 outlines a typical system-level design process. A task graph specification model, used to capture the system's functionality, is introduced in Section 1.2. Section 1.3 describes the individual system design steps using some illustrative examples. Hardware and software synthesis are briefly discussed in Section 1.4. Finally, Section 1.5 gives an overview of this dissertation and highlights the main contributions and achievements of the presented work.

1.1 Embedded System Design Flow

A typical embedded system, as it can be found, for example, in a smart-phone, is shown in Figure 1.1. It consists of heterogeneous components such as software programmable



(a) Embedded architecture: A distributed heterogeneous system

(b) Embedded software

Figure 1.1: Example of a typical embedded system (smart-phone)

processors (CPUs, DSPs) and hardware blocks (FPGAs, ASICs). These components are interconnected through communication links and form a distributed architecture, such

as the one shown in Figure 1.1(a). Analogue-to-digital converters (ADC), digital-to-analogue converters (DAC), as well as input/output ports (I/O) allow the interaction with the environment. A complete embedded system, however, consists additionally of application software (Figure 1.1(b)) that is executed on the underlying hardware architecture (Figure 1.1(a)). Clearly, effective embedded system design demands optimisation in *both* hardware and software parts of the application. When designing an embedded computing system, as part of a new product, it is common to go through several design steps that bring a novel product idea down to its physical realisation. This is usually referred to as system-level design flow. A possible and common design flow is introduced in Figure 1.2. It is characterised by three important design steps: *system specification* (Step A), *co-synthesis* (Step B), as well as concurrent *hardware and software synthesis* (Step C). The remainder of this section briefly outlines this design flow.

Starting from a new product idea, the first step towards a final realisation is *system specification*. At this stage, the functionality of the system is captured using different conceptual models [57] such as natural language, annotated-graphic representations (finite state machines, data-flow graphs), or high-level languages (VHDL, C/C++, SystemC). This design step is indicated as Step A in Figure 1.2. Having specified the system's functionality, the next stage in the design flow is the *co-synthesis*, shown as Step B in Figure 1.2. The goal of co-synthesis is threefold:

Architecture allocation Firstly, an adequate target architecture needs to be allocated, i.e., it is necessary to determine the quantity and the types of different interconnected components that form the distributed embedded system. Components that can be allocated are given in a predefined technology library.

Application mapping Secondly, all parts of the system specification have to be distributed among the allocated components, that is, tasks (function fragments) and communications (data transfers between tasks) are uniquely mapped to processing elements and communication links, respectively.

Activity scheduling Thirdly, a correct execution order of tasks and communications has to be determined, i.e., the activities have to be scheduled under the consideration of interdependencies.

These three co-synthesis stages aim to optimise the design according to objectives set by the designer, such as power consumption, performance, and cost. In order to reduce

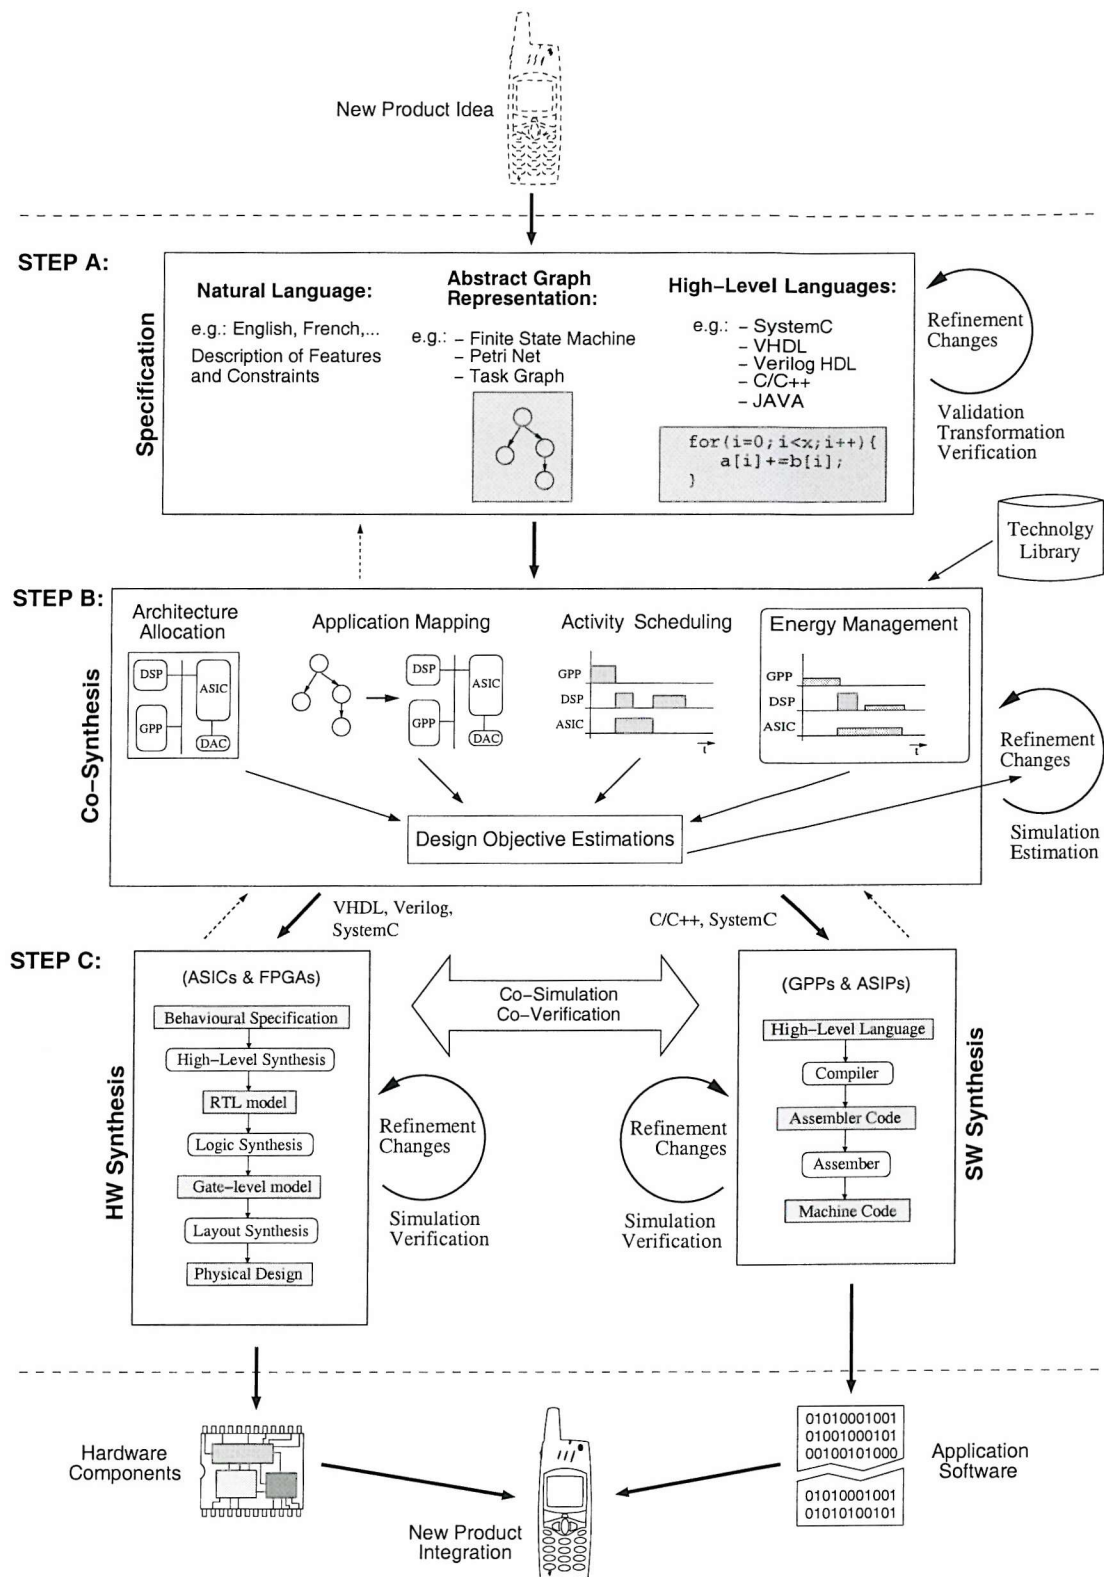


Figure 1.2: Typical design flow of a new embedded computing system

the power consumption, emerging co-synthesis approaches (as the one proposed in this work) tightly integrate the consideration of *energy management techniques* within the design process [63, 72, 89, 90].

Energy management Energy management techniques utilise existing idle times to reduce the power consumption by either shutting down the idle components or by reducing the performance of the components.

The consideration of energy management techniques during the co-synthesis allows the optimisation of allocation, mapping, and scheduling towards their effective exploitation. After the co-synthesis has allocated an architecture as well as mapped and scheduled the system activities (tasks and communications), the next stage in the design flow is the concurrent *hardware and software synthesis*, indicated as Step C in Figure 1.2. These separated design steps transform the system specification, which has been split between hardware and software, into physical implementations. System parts that are mapped onto customised hardware are designed using high-level [8, 18, 56, 128, 146], logic [9, 38, 98, 117], and layout [52] synthesis tools. While system parts that have been mapped onto software programmable processors (CPUs, DSPs) are compiled into assembler and machine code, using either standard or specialised compilers and assemblers [1, 83]. The main advantage of a concurrent hardware (HW) and software (SW) synthesis is the possibility to co-simulate both system parts, with the aim of finding errors in the design as early as possible to avoid expensive re-designs. The following section describes the whole design process shown in Figure 1.2 in more detail and introduces the terminology used throughout this thesis.

1.2 System Specification (Step A)

The functionality of a system can be captured using a variety of conceptual specification models [57]. Different modelling styles are, for example, high-level languages (hardware description and programming languages) such as SystemC, Verilog HDL, VHDL, C/C++, or JAVA, as well as more abstract models such as block diagrams, task graphs, finite state machines (FSMs), Petri nets, or control/dataflow graphs. Typical applications targeted by the presented work can be found in the audio and video processing domain (e.g. multi-media and communication devices with extensive data stream operations).

Such applications fall into the category of data-flow dominated systems. An appropriate representation for these systems is the task graph model [76, 100, 149], which will be introduced in the following section.

1.2.1 Task Graph Representation

The functionality of a complex system with intensive data stream operations can be abstracted as a directed, acyclic graph (DAG) $G_S = (\mathcal{T}, \mathcal{C})$, where the set of nodes $\mathcal{T} = \{\tau_0, \tau_1, \dots, \tau_n\}$ denotes the set of tasks to be executed, and the set of directed edges \mathcal{C} refers to communications between tasks, with $\gamma_{ij} = (\tau_i, \tau_j) \in \mathcal{C}$ indicating a communication from task τ_i to task τ_j . A task can only start its execution after all its ingoing communications have finished. Each task can be annotated with a deadline θ , the time by which its execution has to be finished. Furthermore, the task graph inherits a repetition period ϕ which specifies the maximal delay between two invocations of the source tasks (tasks with no ingoing edges). Structurally, task graphs are similar to the data-flow graphs that are commonly used in high-level synthesis [56, 146]. However, while nodes in data-flow graphs represent single operations, such as multiplications and additions, the nodes in task graphs are associated with larger (coarse) fragments of functionality, such as whole functions and processes. The concept behind this model can be exemplified using a simple illustrative example.

Example 1: For the purpose of this example, consider an MP3 audio decoder. In order to reconstruct the "original" stereo audio signal from an encoded stream, the decoder reads the data stream and applies several transformations such as Huffman decoding, dequantisation, inverse discrete cosine transformation (IDCT), and antialiasing. A possible task graph specification along with a high-level language description in C of such an MP3 decoder is shown in Figure 1.3. The figure outlines the relation between task graph model and high-level description. In this particular example the granularity of each task in the task graph corresponds to a single sub-function of the C specification. For instance, the Huffman Decoder tasks (τ_3 and τ_4) in Figure 1.3(a) reflect the functionality that is performed by the third sub-function in Figure 1.3(b). The flow of data is expressed by edges between the individual tasks. The output data produced by the Huffman Decoder tasks, for example, is the input of the dequant tasks (τ_5 and τ_6), indicated by the communication edges $\gamma_{3,5}$ and $\gamma_{4,6}$. In order to decode the compressed data into a high quality audio

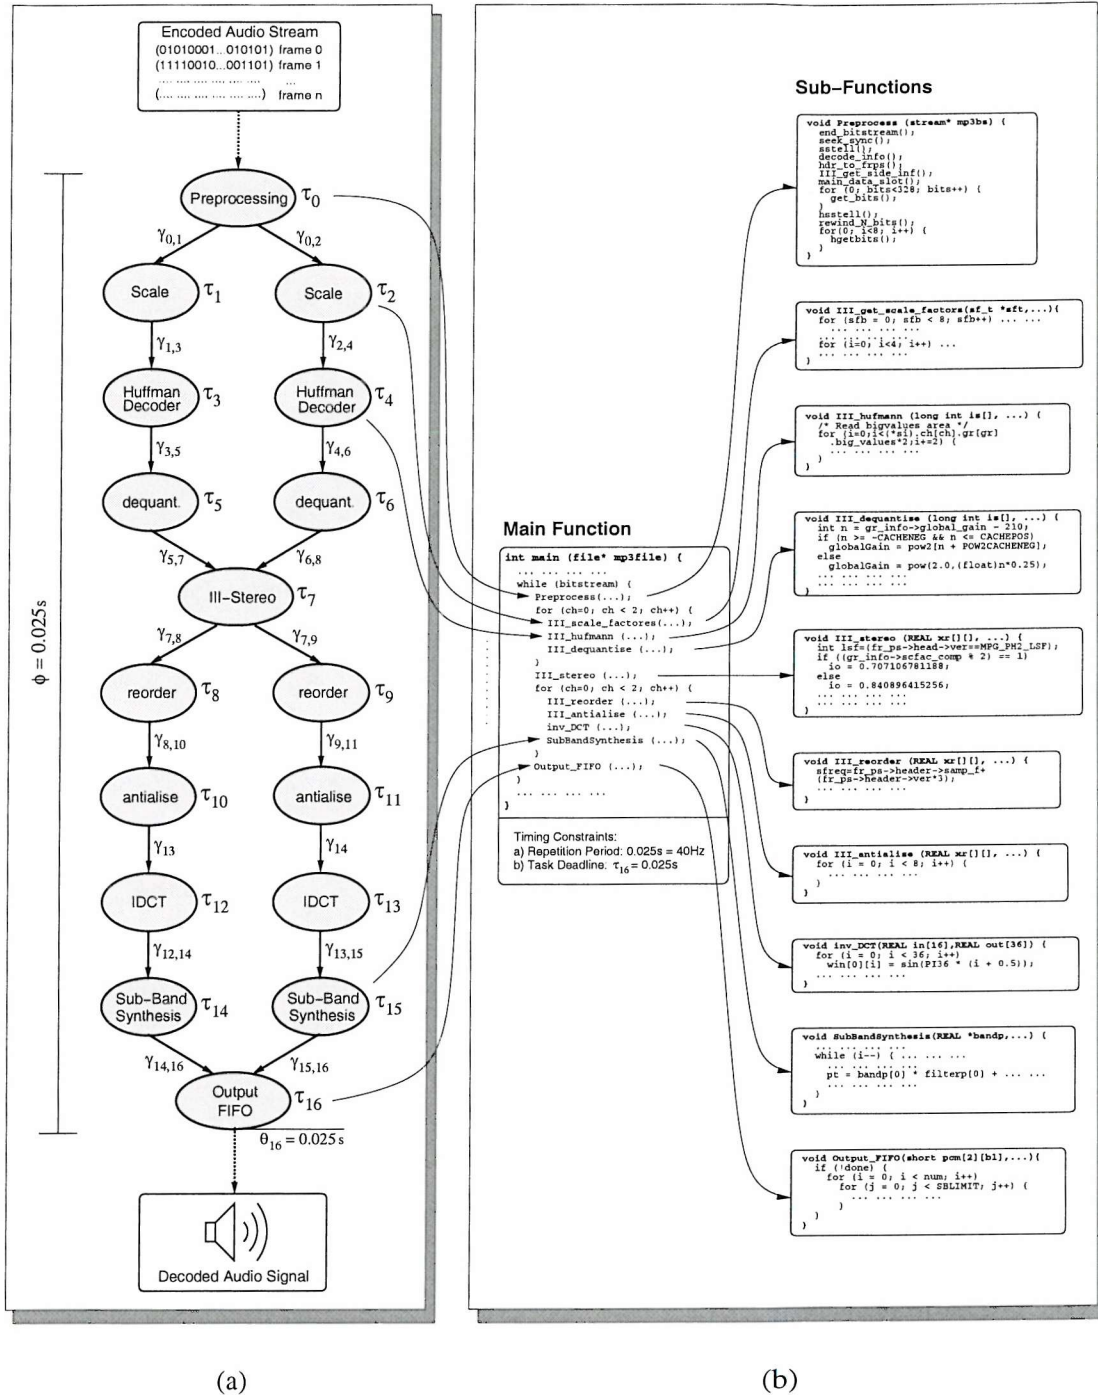


Figure 1.3: MP3 decoder given as (a) task graph specification (17 tasks and 18 communications) and (b) high-level language description in C

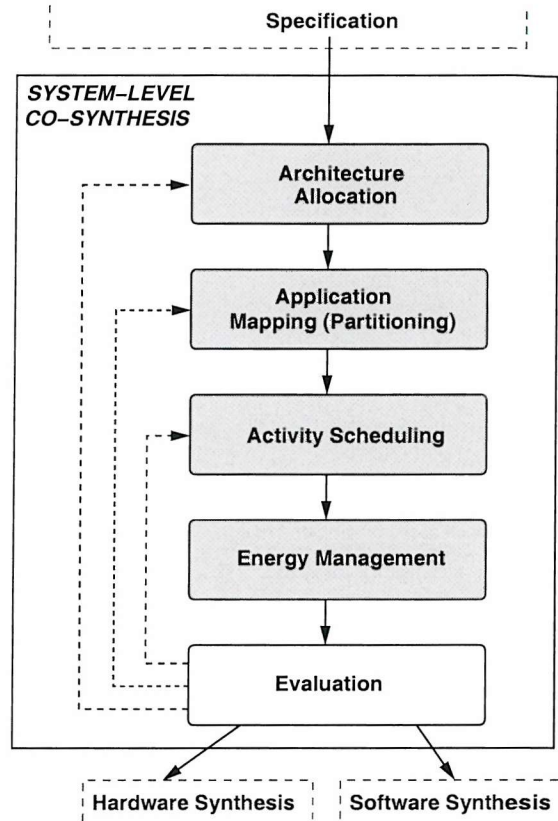


Figure 1.4: System-level co-synthesis flow

signal, one execution of all tasks in the graph, starting from task τ_0 and finishing with τ_{16} , has to be performed in at most $25ms$ as expressed by the task deadline θ_{16} . However, to obtain real-time decompression of a continuous music stream, the execution of all tasks has to be performed 40 times per second, i.e., with a repetition rate of $\phi = 25ms$. Although in this particular example the deadline and the repetition rate are identical, they might vary in other applications. As opposed to the C specification, the task graph explicitly exhibits *application parallelism* as well as *communication between tasks (data flow)*, while the exact algorithmic implementation of each function is abstracted away. \square

1.3 Co-Synthesis (Step B)

Once the system's functionality has been specified as task graph, the system designers will start with the system-level co-synthesis. This is indicated as Step B in Figure 1.2. In addition, Figure 1.4 shows the co-synthesis flow in diagrammatic form. Co-synthesis

is the process of deriving a mixed hardware/software implementation from an abstract functional specification of an embedded system. To achieve this goal, the co-synthesis needs to address four fundamental design problems: *architecture allocation*, *application mapping*, *activity scheduling*, and *energy management*. Figure 1.4 shows the order in which these problems have to be solved. In general, these co-synthesis steps are iteratively repeated until all design constraints and objectives are satisfied [48, 50, 66, 148]. An iterative design process has the advantage that valuable feedback can be provided to the different synthesis steps. This feedback, which is indicated by dashed upwards arrows in the figure, is used to guide the optimisation process towards the satisfaction of design constraints. The following sections explain the co-synthesis flow shown in Figure 1.4 and the four subproblems in more detail.

1.3.1 Architecture Allocation

One of the first questions that needs answering during the design of a new embedded system is what system components (processing elements and communication links) should be used in order to implement the desired product functionality. This part of the co-synthesis is known as *architecture allocation*. Generally, there are many different target architectures that can be used to implement the desired functionality. Problematic, however, is the correct choice as indicated in Figure 1.5. The overall goal of the co-synthesis process is to identify the "most" suitable architecture. Certainly, the "most" suitable architec-




Figure 1.5: Architectural selection problem

ture should provide enough performance for the application in order to satisfy the timing constraints, while, at the same time, cost, design time, and energy dissipation should be

reduced to a minimum. The importance of architecture allocation becomes clearer when considering the advantages and disadvantages associated with processing elements of various kinds. Table 1.1 gives the most relevant component trade-offs. Consider, for instance,

	GPPs	ASIPs	FPGAs	ASICs
Cost	++	++	o	--
Flexibility	++	+	o	--
Performance	-	o	+	++
Energy-Efficiency	-	o	+	++

Table 1.1: Trade-offs between several heterogeneous components
 (+ + \equiv highly advantageous, + \equiv advantageous, o \equiv moderate,
 - \equiv disadvantageous, - - \equiv highly disadvantageous)

the two processing elements (PEs): general-purpose processor (GPP) and application specific integrated circuit (ASIC). While software implementations on off-the-shelf GPPs are more flexible and cheaper to realise than hardware designs, the ASIC offers higher performance and better energy-efficiency. Similarly, the application specific instruction set processors (ASIPs) and field-programmable gate arrays (FPGAs) show different trade-offs. Of course, the non-recurring engineering cost (NRE) is mainly important for low volume products. For high volume applications this cost is amortised and becomes less important. Certainly, selecting the appropriate system components, in order to balance between these trade-offs, is of utmost importance for high quality designs. The intention of system-level co-synthesis tools is to aid the system designer in effectively exploring the architectural design space, in order to find a suitable target architecture rapidly.

1.3.2 Application Mapping

Following the co-synthesis flow given in Figure 1.4, the next step after architecture allocation is *application mapping*. During this step the tasks and communications of the system specification are mapped onto the allocated processing elements (PEs) and communication links (CLs) of the architecture, respectively. Figure 1.6 illustrates two different mappings of a system specification onto identical target architectures. These two mappings differ in the assignment of task τ_4 , which is either mapped to the ASIC (Mapping 1) or to CPU2 (Mapping 2). Mapping explicitly determines if a task is implemented in hardware or software, hence, the term *hardware/software partitioning* is often mentioned in this context. Due to the heterogeneity of processing elements, the mapping specifies the

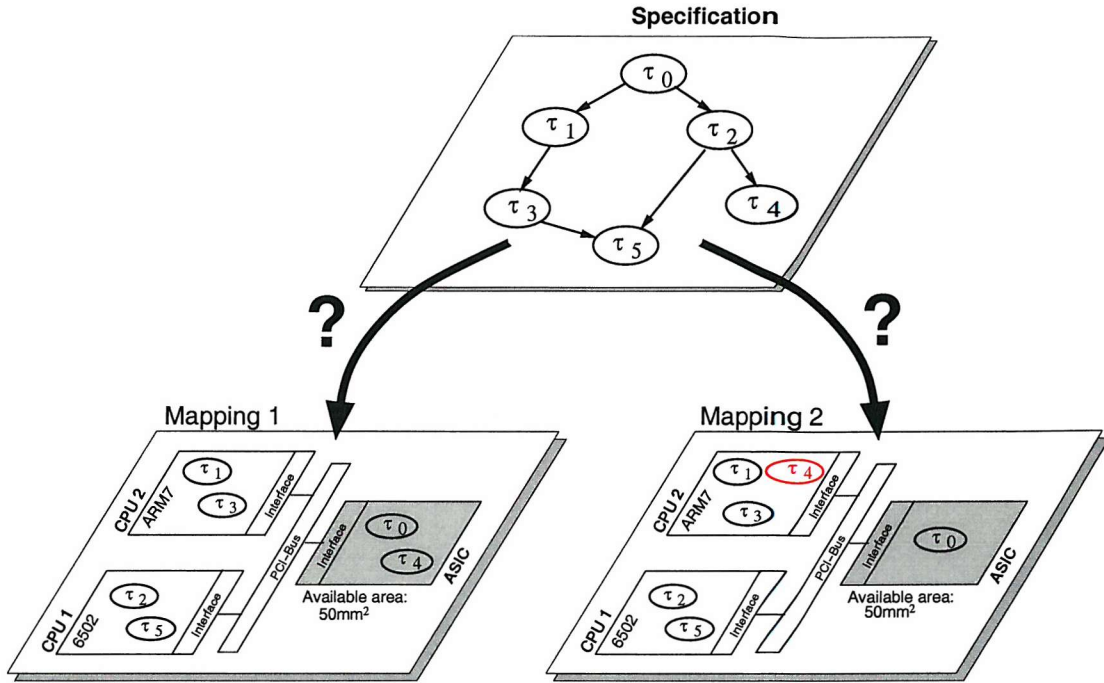


Figure 1.6: Application mapping onto hardware and software components

execution characteristics of each task and communication. Consider, for example, the execution characteristics of the tasks shown in Table 1.2. This table gives the execution times

Task	CPU 1 (6502 @10MHz)		CPU 2 (ARM7 @20MHz)		ASIC (50mm ² , Technology: 0.6μm)		
	t_{exe} (ms)	P_{dyn} (mW)	t_{exe} (ms)	P_{dyn} (mW)	t_{exe} (ms)	P_{dyn} (mW)	A (mm ²)
τ_0	89.3	3.6	12.1	23	1.8	0.13	7.76
τ_1	25.2	3.9	3.0	26	0.3	0.05	5.82
τ_2	19.7	4.4	2.8	28	0.2	0.07	9.71
τ_3	31.1	3.8	4.7	28	0.4	0.02	12.52
τ_4	172.2	3.9	22.3	27	2.7	0.12	8.05
τ_5	27.2	4.2	3.5	24	0.6	0.02	3.74

Table 1.2: Task execution properties (time and power) on different processing elements

t_{exe} and power dissipations P_{dyn} of each task in the specification of Figure 1.6, depending on the mapping to a 6052 8-bit microprocessor (running at 10MHz), an ARM7TDMI 32-bit microprocessor (running at 20MHz), or an ASIC in 0.6μm technology which offers a usable die size of 50mm². In addition to the time and power values, the hardware area A required for tasks implemented on the ASIC is given. In general, hardware implementations are more efficient in terms of performance and power consumption than software

realisations. However, the design of hardware is a more time consuming process. Clearly, determining a good mapping solution is of crucial importance for the system design. Inappropriately distributing the activities among the components can result in poor utilisation of the system, necessitating the allocation of an architecture with higher performance, hence, increasing the system cost.

1.3.3 Activity Scheduling

Moving further in the design flow of Figure 1.4, the next step after application mapping is *activity scheduling*. The function of scheduling is to order the execution of tasks and communications (both activities) such that timing constraints are satisfied. This is not a trivial problem, since several activities mapped onto the same component cause congestion, which, in turn, hampers the effective exploitation of parallelism in the application. Hence, a good schedule should allow to exploit this parallelism effectively in order to improve the system performance.

Given an allocated architecture and a mapping of tasks and communication as well as a task graph specification, Figure 1.7 depicts two possible schedule solutions (Schedule 1 and Schedule 2). According to the system specification, the execution of the tasks τ_4 and τ_5 must be finished before deadline θ is exceeded. Thus, if the deadlines are violated the schedule is invalid. Consider the following scheduling scenarios given in Schedule 1 and Schedule 2 of Figure 1.7. After the initial task τ_0 has finished its execution, the communications $\gamma_{0,1}$ and $\gamma_{0,2}$ become ready. However, since both communications need to share the same bus it is necessary to sequence the transfers, since only one transfer is possible at a given time. Thus, a scheduling decision has to be taken at this point. The first schedule shown in Figure 1.7 corresponds to a schedule in which communication $\gamma_{0,1}$ takes place before communication $\gamma_{0,2}$. As it can be observed from this schedule, the executions of tasks τ_4 and τ_5 finish before deadline θ , hence, this solution represents a valid schedule ($t_{f1} < \theta$). On the other hand, if communication $\gamma_{0,2}$ is scheduled before communication $\gamma_{0,1}$, as shown in Schedule 2, the execution of task τ_1 is delayed, which further delays task τ_3 and communication $\gamma_{3,5}$. Ultimately, the execution of task τ_5 starts too late to finish the execution before deadline θ . Thus, the second schedule represents an invalid solution ($t_{f2} > \theta$).

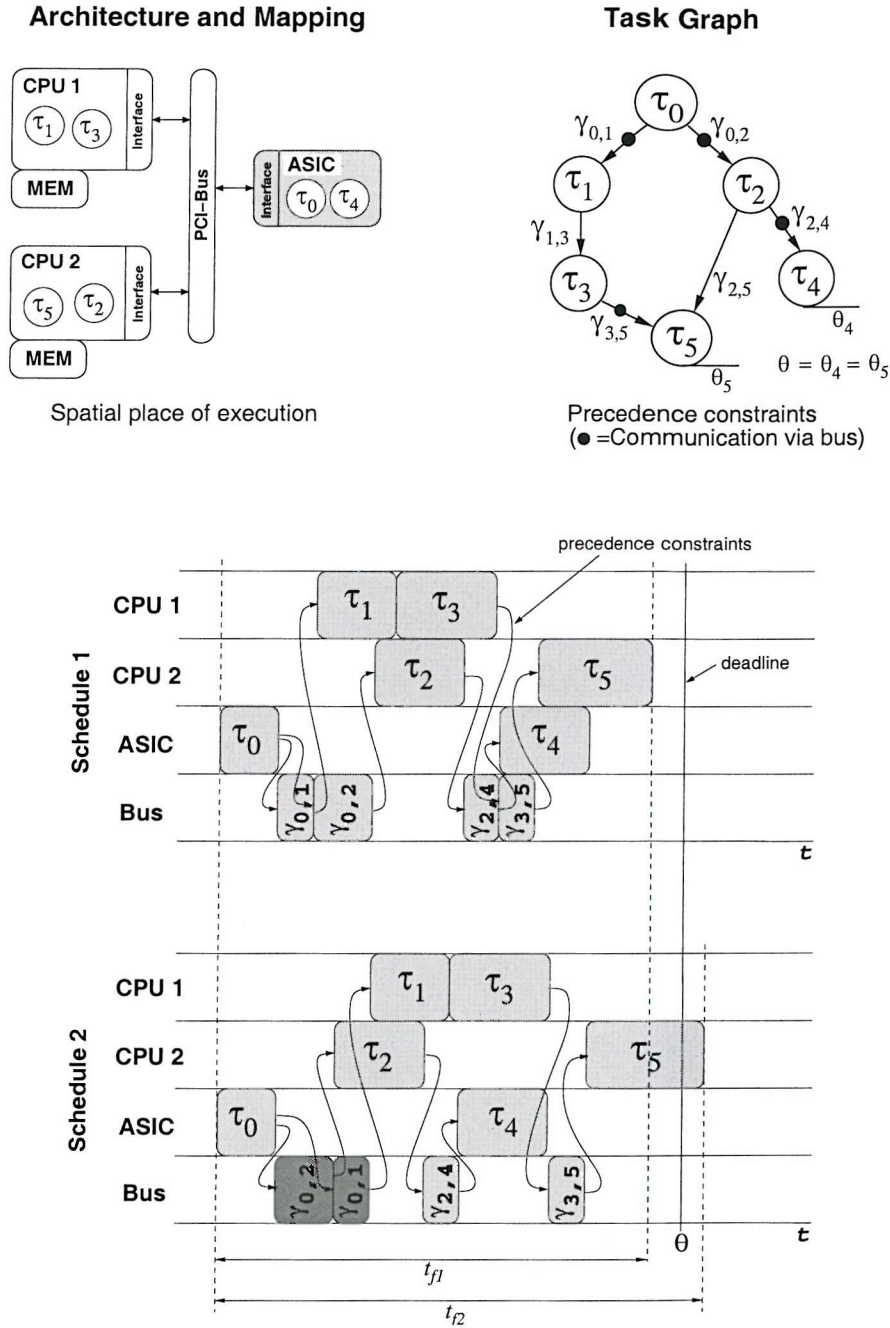


Figure 1.7: Two different scheduling variants based on the same allocated architecture and identical application mapping

1.3.4 Energy Management

Having allocated an architecture as well as having mapped and scheduled the application onto it, the next step within the co-synthesis flow of Figure 1.4 is the utilisation of *energy management techniques*. This step is necessary to accurately estimate the energy require-




Figure 1.8: System schedule with idle and slack times

ments of the system, which is used to guide the optimisation of allocation, mapping, and scheduling towards energy-efficient designs. In general, energy management techniques exploit idle times and slack times within the system schedule by shutting down processing elements (PEs) [25, 87] or by reducing the performance of individual PEs [33, 144]. Idle times and slack times are defined as follows:

- **Idle times** refer to periods in the schedule when PEs and CLs do not experience any workload, i.e., during these intervals the components are redundant (see Figure 1.8).
- **Slack times** is the difference between task deadline and task finishing time of sink tasks (tasks with no outgoing edges), i.e., slack times are a result of over-performance (see Figure 1.8). Clearly, slack time is a special case of idle time.

Two important energy management techniques are dynamic power management (DPM) [25, 87, 134] and dynamic voltage scaling (DVS) [33, 72, 74, 144]. DPM puts processing elements and communication links (both components) into standby or sleeping modes whenever they are idle. Nevertheless, the reactivation of components takes finite time and energy; hence, components should only be switched off or set into a standby mode if the idle periods are long enough to avoid deadline violations or increased power consumptions [25, 88]. DVS, on the other hand, exploits slack time by reducing simultaneously clock frequency and supply voltage of PEs. Thereby, DVS adapts the component performance to the actual requirement of the system. In this way, substantial saving are achieved since the energy consumption of the system components is proportional to the square of the supply voltage ($E \propto V^2$) [35]. The basic concept behind DVS is demonstrated in Figure 1.9. It can be observed from Figure 1.9(a) that tasks τ_4 and τ_5 finish execution before deadline θ . As indicated in the figure, this results in slack time. Instead of switching-off

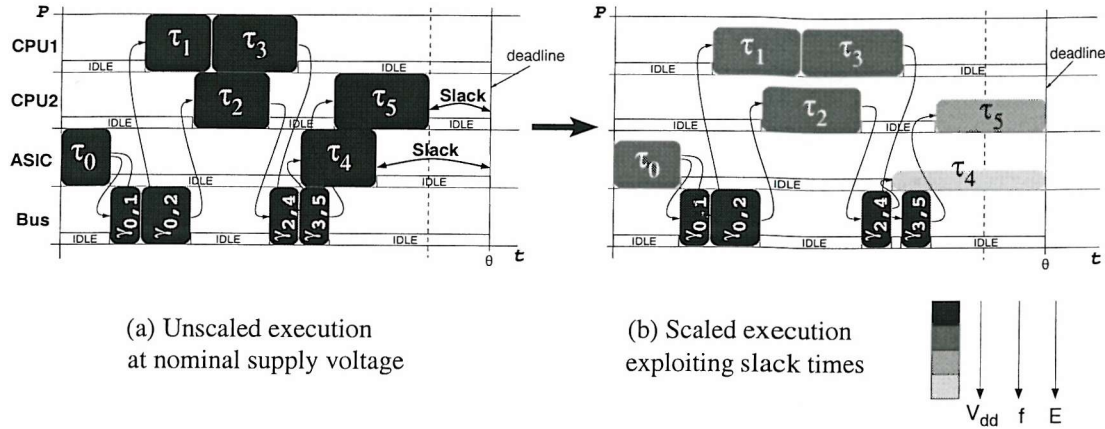


Figure 1.9: The concept of dynamic voltage scaling

the components during these times (as done by DPM), it is possible to prolong the execution of all six tasks. This is achieved by scaling down the supply voltage and frequency of the processing elements until the tasks τ_4 and τ_5 just finish on time (as shown in Figure 1.9(b)). The main problem that needs to be addressed here is how to distribute the available slack time among the tasks, in order to achieve the "highest" possible energy savings.

Nevertheless, the effectiveness with which DPM and DVS can be applied depends significantly on the available idle and slack times. A worthwhile optimisation of allocation, mapping, and scheduling must take the optimisation of idle and slack time into account, in order to allow a most effective exploitation of both techniques [64, 72, 88, 89]. In general, such an optimisation requires the iterative execution of the co-synthesis steps (allocation, mapping, scheduling), until the "most" suitable implementation of the system has been found [63, 89].

1.4 Hardware and Software Synthesis (Step C)

The previous section has outlined the system-level co-synthesis (Step B in Figure 1.2), which transforms an abstract specification into an architectural description of a mixed hardware/software system. The final step in the embedded system design flow is the concurrent *hardware and software synthesis* (Step C in Figure 1.2). This step brings the mixed hardware/software description of the system down to a physical implementation, i.e., the specification fragments (tasks) that have been distributed among the hardware

and software components of the system need to be realised. This is achieved through two separate, yet concurrent synthesis steps: hardware synthesis and software synthesis. One of the main advantages of concurrent HW/SW design is the ability to check the correctness of the overall system by means of simulation, i.e., the interaction between hardware and software can be co-simulated [115]. Note, whereas system-level co-synthesis targets the design of interacting components, the main aim of hardware and software synthesis is the design of the *individual* hardware components and the software tasks running on programmable processors.

Hardware Synthesis: The design of complex hardware components is based on existing very large scale integration (VLSI) synthesis tools [8, 9, 39, 128, 146, 147]. Figure 1.10 illustrates a possible hardware synthesis process that consists of three subsequent design steps.

- (a) A *high-level synthesis tool* (or behavioural synthesis tool) [8, 128, 146] transforms a behavioural specification into a structural description at the register-transfer level (RTL). Here the individual components are represented by data paths which execute arithmetic operations under control of a control unit.
- (b) The RTL description (e.g. in structural VHDL) is then translated into a gate-level representation using a *logic synthesis tool* [9, 10]. In this stage of the design, the control unit as well as the data path are structurally represented as netlists of logic gates.
- (c) The final layout mask (used for IC fabrication) is generated from the gate-level description through a *layout synthesis tool* [52]. Here the individual physical gates are placed and interconnections are routed.

It should be noted that power reduction can be addressed at all three synthesis stages (high-level: e.g. clock-gating [27, 147], gate-level: e.g. logic optimisation [41, 95], mask-level: e.g. technology choice [34, 41]). However, independent of these low-level power reduction techniques, the previously discussed energy management techniques (DPM and DVS) can be applied at a higher level of abstraction (system-level) to further improve the savings in energy. In general, the higher the level of abstraction at which the energy minimisation is addressed, the higher are the achievable energy saving [110].

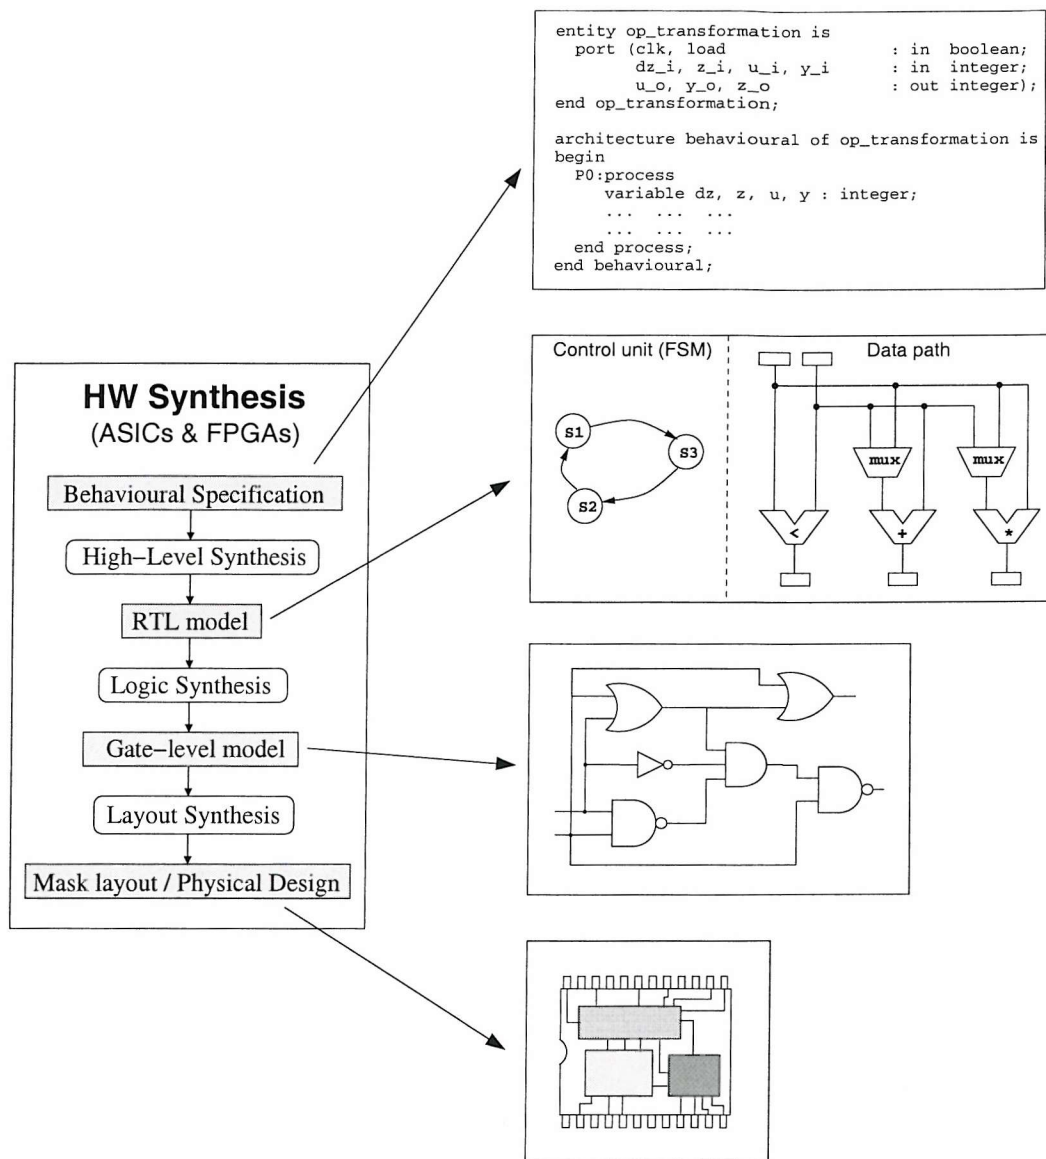


Figure 1.10: Hardware synthesis flow

Software Synthesis: Similarly to the hardware synthesis, all tasks that have been mapped to software programmable components have to be transformed from a high-level description (e.g. C/C++, JAVA, SystemC) into low-level machine code. A software translation hierarchy is shown in Figure 1.11 and consists of two steps:

- (a) The initial specification in a high-level language is *compiled* into assembly code. This is carried out either using standard compilers, such as GCC [1, 2], or using specialised compilers that are optimised towards specific processor types (e.g. DSPs) [83]. The goal of the optimisation is the effective assignment of variables to reg-

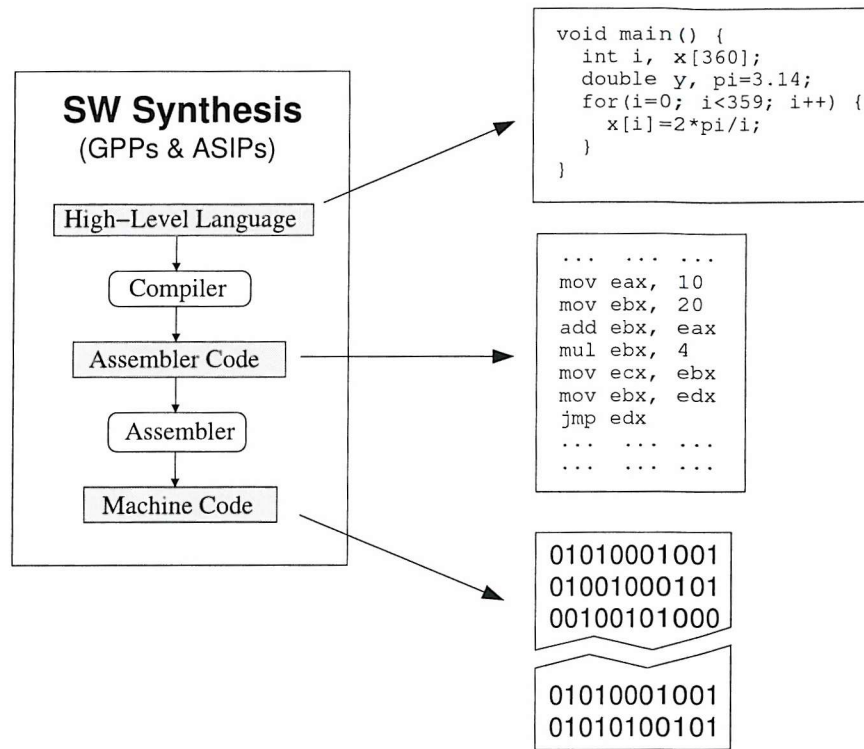


Figure 1.11: Software synthesis flow

isters such that operations can be performed without "time consuming" memory accesses.

- (b) Once an optimised assembly code has been generated, the low-level code generation is carried out by processor specific assemblers that translate the assembler code into executable machine code.

There exist also techniques for compiler-based power minimisation such as instruction reordering and reduction of memory accesses [84, 139, 140]. Further, sizeable power saving can be obtained through a careful algorithmic design at the source code level [133]. Clearly, such software power minimisation approaches and system-level energy management techniques do not exclude each other. In fact, for a most energy-efficient system design both techniques should be considered.

1.5 Contributions and Thesis Overview

This work presents novel techniques and algorithms for the automated design of energy-efficient distributed embedded system. In particular, the energy reduction capabilities of dynamic voltage scaling (DVS) are investigated and analysed in the context of highly programmable embedded systems with strict performance and cost requirements. The remainder of this thesis is organised as follows. Chapter 2 provides a survey of the most relevant and related works and outlines the necessary background information that is helpful for the understanding of the discussed subject.

Chapter 3 introduces a novel technique for dynamic voltage scaling for distributed architectures that effectively reduces the energy dissipation of the embedded system. This technique addresses the energy management problem discussed in Section 1.3. The proposed approach considers the power variations inherent to the execution of different tasks, in order to increase the efficiency with which DVS can be applied.

Based on this DVS technique, Chapter 4 introduces a new co-synthesis approach for distributed embedded system that potentially contain voltage-scalable components. Application mapping and activity scheduling are optimised towards the effective utilisation of DVS, i.e., towards energy reduction. This optimisation simultaneously aims at the identification of solution candidates that fulfil the imposed timing constraint and reduced the system cost.

Chapter 5 further extends the proposed co-synthesis approach towards the design of multi-mode embedded systems which integrate several different applications into a single device. The introduced multi-mode co-synthesis aims at energy-efficiency as well as cost effective utilisation of the hardware components. It is demonstrated that substantial energy savings can be achieved without modification of the underlying hardware architecture, even when neglecting DVS.

The techniques introduced in the preceding chapters and their algorithmic implementations have been combined into a new prototype co-synthesis tool for energy-efficient embedded systems. This tool is introduced in Chapter 6 and its usage is demonstrated using a real-life smart-phone that merges a cellular GSM phone, a digital camera, and an MP3-player into one device. Chapter 7 concludes the presented work and outlines potential areas of future research.

List of Publications

The work presented in this thesis has resulted in a number of original conference and journal publications.

Presented Workshop and Conference Papers:

- *"Energy Minimisation for Processor Cores using Variable Supply Voltages"*, Marcus T. Schmitz and Bashir M. Al-Hashimi, IEE System on a Chip Workshop, September 2000, Cork, Ireland. [120]
- *"Low Power Process Assignment for Distributed Embedded Systems Using Dynamic Voltage Scaling"*, Marcus T. Schmitz and Bashir M. Al-Hashimi, IEE Hardware-Software Co-Design, December 2000, London, United Kingdom. [121]
- *"Considering Power Variations of DVS Processing Elements for Energy Minimisation in Distributed Systems"*, Marcus T. Schmitz and Bashir M. Al-Hashimi, 14th International Symposium on System Synthesis (ISSS 2001), pp. 250-255, October 2001, Montreal, Canada. [122]
- *"Energy-Efficient Mapping and Scheduling for DVS Enabled Distributed Embedded Systems"*, Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles, Design, Automation and Test Conference 2002 (DATE 2002), pp. 514–521, March 2002, Paris, France. [123]
- *"Synthesising Low Power Distributed Systems: A Multi-Mode System Approach"*, Marcus T. Schmitz and Bashir M. Al-Hashimi, 2nd UK ACM SIGDA Workshop 2003, 16–17 September 2002. [119]
- *"A Co-Design Methodology for Energy-Efficient Multi-Mode Embedded Systems with Consideration of Mode Execution Probabilities"*, Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles. Accepted for publication in the proceedings of the Design, Automation and Test Conference 2003 (DATE 2003) pp. 960–965, March 2003, Munich, Germany. [125]

Journal Papers:

- *"Synthesising Energy-Efficient Embedded Systems with LOPOCOS"*, Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles, Design Automation of Embedded Systems, Volume 6:4, pages 401–424, Kluwer Academic Publishers, July 2002. [124]

- *"Iterative Schedule Optimisation for Voltage Scalable Distributed Embedded Systems"*, Marcus T. Schmitz, Bashir M. Al-Hashimi, Petru Eles. Accepted for publication in ACM Transactions on Embedded Computing Systems. [126]
- *"System-Level Design of Energy-Efficient Multi-Mode Embedded Systems"*, Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles. Currently under preparation for submission to IEEE Transactions on CAD. [127]

Chapter 2

Background and Previous Work

Reducing power consumption has emerged as a primary design goal, in particular for battery-powered embedded systems. Low power design techniques for digital components have been intensively investigated over the last decade [26, 96, 103, 109, 141, 147]. These techniques focus mainly on the optimisation of a single hardware component in isolation. However, embedded systems are often far more complex than single components — they consist of several interacting heterogeneous components. Here the interrelation between the different processors and hardware blocks should be carefully considered during the synthesis in order to achieve an energy-efficient design. Two techniques that can be used for energy minimisation of distributed embedded systems are: dynamic power management (DPM) [22] and dynamic voltage scaling (DVS) [74, 144]. These system-level energy management techniques achieve energy reductions by selectively switching off unused components (DPM) or by scaling down the performance of individual components in accordance to temporal performance requirements of the application (DVS).

The aim of this chapter is to introduce the sources of power dissipation within distributed embedded systems and to outline how energy management techniques can be applied to reduce the dissipated energy (Section 2.1–2.3). Furthermore, an overview of the most relevant previous work is given, differentiating between general co-synthesis approaches without energy minimisation and co-synthesis approaches with energy minimisation (Section 2.4).

2.1 Energy Dissipation of Processing Elements

The power P_{PE} dissipated by computational components (CPUs, ASIPs, FPGAs, ASICs) of an embedded system, i.e. processing elements, is caused by two distinctive effects. First, *static* currents that occur whenever the processing element is switched on, even when no computations are carried out on this unit. Second, active computations cause switching activity within the circuitry that results in *dynamic* power dissipation whenever computations are performed. Accordingly to both sources the total power dissipation of processing elements is given by:

$$P_{PE} = P_{static} + P_{dynamic} \quad (2.1)$$

Both static and dynamic power dissipations can be further subdivided into two power components, each [34]:

$$P_{PE} = \underbrace{P_{leak} + P_{bias}}_{P_{static}} + \underbrace{P_{sc} + P_{sw}}_{P_{dynamic}} \quad (2.2)$$

The static power P_{static} has two parts: leakage power P_{leak} and bias power P_{bias} . While the dynamic power $P_{dynamic}$ consists of short-circuit power P_{sc} and switching power P_{sw} . Out of these four source of power dissipation, switching power P_{sw} is the dominant one which accounts for approximately 90% of the total PE power consumption [35]. The following discussion concentrates on this portion of the total power dissipation, simply referred to as power or dynamic power. It should be noted, however, that with shrinking feature size ($< 0.07\mu m$) and reduced threshold voltage levels, the leakage currents become additionally an important issue [34, 116].

Switching power P_{sw} is dissipated due to the charging and discharging of the effective circuit load capacitance C_L (parasitic capacitors of the circuit gates). To clarify the source of switching power P_{sw} , consider the simple gate-level circuit shown in Figure 2.1(a) and in particular the inverter gate shown in Figure 2.1(b). This inverter undergoes the following transitions. First, the input signal y is set to high (1), i.e., Tr1 is open (not conducting) while Tr2 is conducting. Accordingly, the circuit load capacitance C_L is discharged since Tr2 pulls the capacitance to ground. The load capacitance C_L represents the intrinsic capacitance of the inputs v and w of the AND and NOR gates. Now consider a transition from high (1) to low (0) at the input of the inverter y . In this case the transistor Tr2 is open and

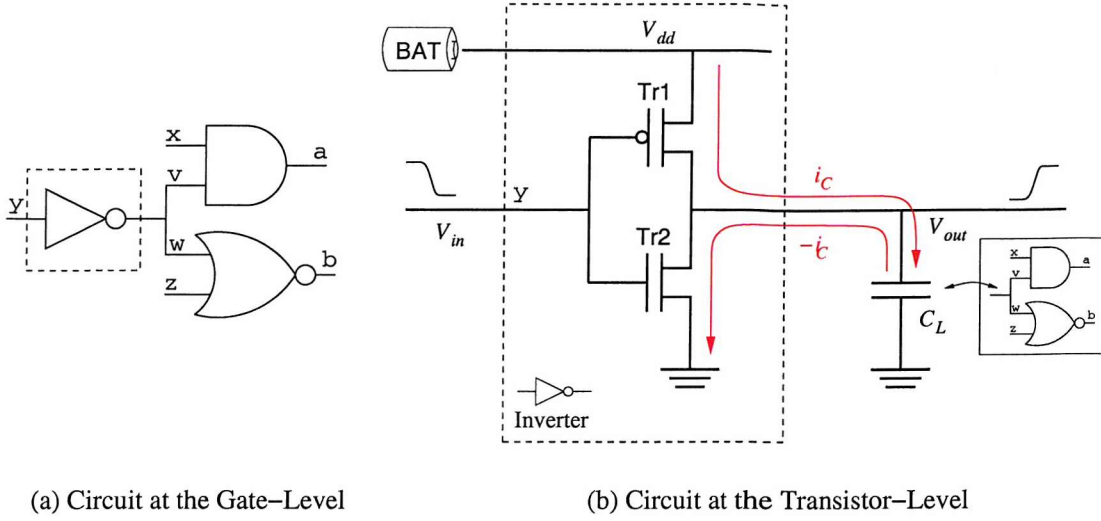


Figure 2.1: Dynamic power dissipation of an inverter circuit [34]

Tr1 connects the capacitance C_L to the supply voltage V_{dd} source, charging C_L via i_C . The power dissipated by this transition is given by:

$$P_{sw} = V_{dd} \cdot i_C \quad (2.3)$$

where the dynamic current i_C changes according the dynamic voltage on the output:

$$i_C = C_L \cdot \frac{\partial V_{out}}{\partial t} \quad (2.4)$$

Therefore, energy is transferred from the power supply to the load capacitance C_L . However, it can be observed that a transition from low to high at the input does not draw any current from the source, but instead discharges the load capacitance C_L via Tr2. This indicates that power, from the battery point of view, is only dissipated during output transitions from 0 to 1, i.e., when the load capacitance is charged. According to the above given observation, the energy consumption of the circuit is solely caused by transitions from low to high at the output of the gate. The dissipated switching energy¹ of one clock cycle, which takes a time of T , can be calculated as [34]:

$$E_{sw}^{0 \rightarrow 1} = \int_0^T P_{sw} dt = V_{dd} \cdot \int_0^T i_C dt = C_L \cdot V_{dd}^2 \quad (2.5)$$

¹In order to prevent a confound usage of the terms *power* and *energy*, the following definition is used throughout this thesis. The term *power dissipation* refers to the physical value power, while the terms *power consumption*, *energy dissipation*, and *energy consumption* refer to physical value energy.

where the time $T = 1/f$ (period of on clock cycle) depends on the operational frequency f at which the circuit is clocked.

Although the above considerations were restricted to a single inverter gate, the same observations hold for more complex circuits, such as microprocessors [33]. As a result, the total energy E_{dyn}^τ drawn from the batteries by a PE performing a computational task τ depends additionally on the number of clock cycles N_C needed to execute this task and the switching activity α . Therefore, the total energy E_{dyn}^τ is given by:

$$E_{sw}^\tau = N_C^\tau \cdot \alpha \cdot C_L \cdot V_{dd}^2 \quad (2.6)$$

Dividing Equation (2.6) by the execution time of the task ($T \cdot N_C^\tau$), the well known equation for power dissipation due to switching can be derived [34]:

$$P_{sw}^\tau = E_{sw}^\tau \cdot \frac{1}{T \cdot N_C} = \alpha \cdot C_L \cdot V_{dd}^2 \cdot f \quad (2.7)$$

Considering Equations (2.6) and (2.7) leads to very interesting conclusions. If we assume that the load capacitance C_L is a constant given by the complexity of the design and the circuit technology, and further the switching activity is a constant depending on the computational task then:

Although decreasing the operational frequency f leads to a reduction in power dissipation (Equation (2.7)), it does not reduce the energy dissipation (Equation (2.6)), which is important for battery-lifetime. A simple example can be used to illustrate this. Consider a computation that requires 10ms on a PE running at 10MHz and dissipating 200mW. This computation requires an energy of $200\text{mW} \cdot 10\text{ms} = 2\text{mJ}$. If the operational frequency is reduced from 10MHz to 5MHz, the power consumption of the processor decreases to 100mW, according to Equation (2.7). Nevertheless, the frequency reduction increases the computational time from 10ms to 20ms. Therefore, the dissipated energy remains unchanged $100\text{mW} \cdot 20\text{ms} = 2\text{mJ}$. Hence, the only possibility to reduce the energy consumption is to reduce the circuit supply voltage V_{dd} . Of course, reducing the supply voltage necessitates the reduction of the frequency in order to ensure correct operation, as it will be shown in the following.




Figure 2.2: Supply voltage dependent circuit delay

When reducing the supply voltage of a digital circuit, the time required for gate signals to settle is prolonged, which, in turn, increases the circuit delay [34]. The source of this increased circuit delay is simplified shown in Figure 2.2. The two sub-figures show the gate output voltage V_{out} over time during a transition from low to high. While the Figure 2.2(a) corresponds to a transition at maximal supply voltage, Figure 2.2(b) shows a transition at reduced voltage. Subsequent gates recognise the signal as high as soon as the output voltage V_{out} exceeds the threshold voltage V_t . It can be seen that the reduction of the supply voltage V_{dd} leads to longer charging times until the threshold is reached (compare d_1 with d_2).

The circuit delay d , which is inverse proportional to the operational frequency f at which the component is clocked, can be approximated (error < 10%) as [31, 34]:

$$d \propto \frac{1}{f} \approx k_d \cdot \frac{V_{dd}}{(V_{dd} - V_t)^2} \quad (2.8)$$

with the technology dependent constant k_d which is given by:

$$k_d = \frac{C_L}{2 \cdot W \cdot v_{sat} \cdot C_{OX}} \quad (2.9)$$

where the constants W , v_{sat} , and C_{OX} denote width of the sub-micron CMOS device, velocity saturation, and gate capacitance, respectively, which limit the charging current of load capacitance C_L . Based on the energy equation (2.6) and the circuit delay equation (2.8), the normalised energy/delay trade-off can be derived as (see Appendix D):

$$E^*(d^*) = \frac{E_{max}}{V_{max}^2} \cdot \left(V_t + \frac{V_0}{2d^*} + \sqrt{\left(V_t + \frac{V_0}{2d^*} \right)^2 - V_t^2} \right)^2 \quad (2.10)$$

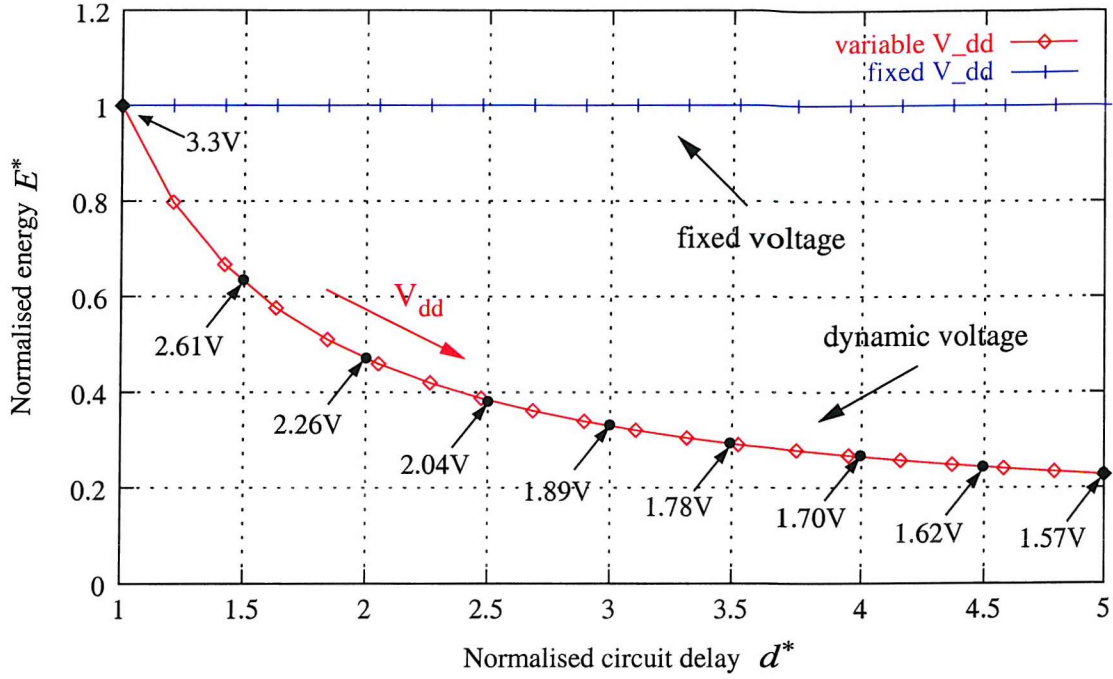


Figure 2.3: Energy versus delay function using fixed and dynamic supply voltages (considering $V_{max} = 3.3V$ and $V_t = 0.8V$)

where E_{max} and V_{max} denote the nominal values of energy dissipation and supply voltage, respectively. The normalised delay $d^* = d(V_{dd})/d(V_{max})$ is represented by the circuit delay at the scaled voltage $d(V_{dd})$ over the minimal circuit delay at nominal voltage $d(V_{max})$. While the constant V_0 is given by:

$$V_0 = \frac{(V_{max} - V_t)^2}{V_{max}} \quad (2.11)$$

Please note that Equation (2.10) differs from the energy/performance trade-off equation given in [67]. The reason for this can be found in the fact that the considered PEs employ lower level power reduction techniques, e.g., gated clocks that avoid switching in unused circuit parts. Based on Equation (2.10), the energy-delay trade-off curve given in Figure 2.3 shows the normalised energy dissipation dependent on the normalised circuit delay for two cases: (a) keeping the supply voltage fixed and (b) dynamically adjusting the supply voltage. For instance, executing a task at 10MHz instead of 30MHz reduces the energy dissipation to approximately 33% of the nominal value when the supply voltage is lowered accordingly. On the other hand, if the supply voltage is kept fixed, the energy dissipation remains constant at the nominal value (see Equation (2.6)).

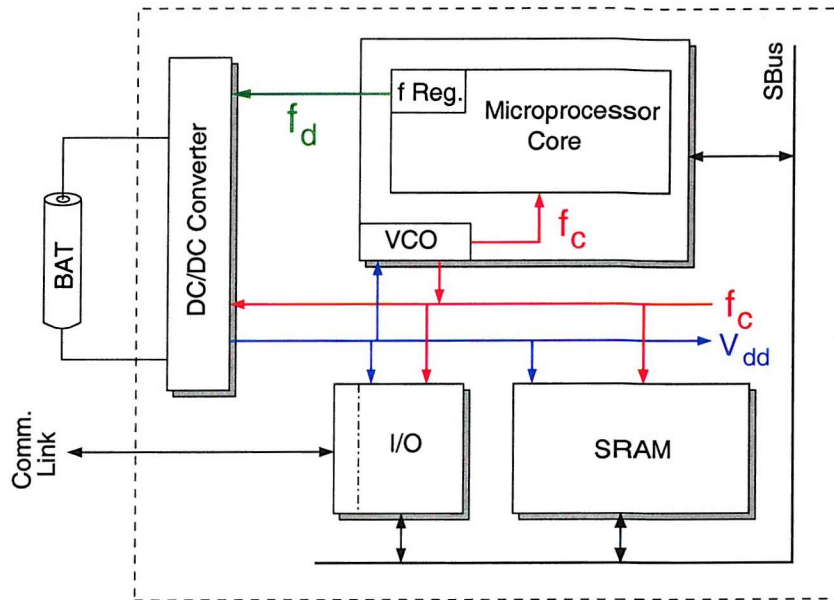


Figure 2.4: Block diagram of DVS-enabled processor [33]

2.2 Energy Minimisation Techniques

Having introduced the fundamentals of power dissipation in digital circuitry, this section outlines two energy reduction techniques that have received considerable attention from academia and industry: dynamic voltage scaling and dynamic power management.

2.2.1 Dynamic Voltage Scaling

A relatively new energy minimisation technique, which leverages the energy-delay trade-off described above (Section 2.1), is dynamic voltage scaling (DVS) [74, 144, 152]. DVS-enabled processors have the ability to dynamically change their supply voltage and operational frequency settings during run-time of the application. Hence, temporal performance requirements of applications can exploit the energy/delay trade-off to reduce energy dissipation. Figure 2.4 shows a block diagram of a typical DVS-enabled processor [33]. A microprocessor core carries out the required computations. This processing unit is connected through a system bus to the static memory (cache unit) and the I/O bus interface. The heart of this system, which enables a dynamic voltage selection, consists of a DC/DC voltage converter, a specialised frequency register, and a voltage controlled oscillator (VCO). Supply voltage and operational frequency are changed by writing the desired frequency f_d into the frequency register, i.e., these changes are carried out un-

der software control. Upon writing the desired frequency into the register, the DC/DC converter compares this frequency with the current frequency f_c (which clocks the microprocessor core, cache, and I/O interface) and either increases or decreases the supply voltage V_{dd} . According to the changed voltage, the VCO adapts the system clock to a higher or lower frequency f_c . Certainly, the whole voltage scaling process requires a finite time. Typical transition times are in the range of tenths of microseconds. For instance, the prototype processor introduced in [33] is able to switch from 5MHz (1.2V) to 80MHz (3.8V), which represents a transition over the full voltage range, in around $70\mu\text{s}$. Nevertheless, the execution of instructions can be seamlessly continued during this transition.

Various chip makers have recently introduced processors with DVS capability. For example, Transmeta introduced the Crusoe processor in 2000. This processor uses a DVS technique called *LongRun*, which enables the TM5600 model to run between 300MHz (1.2V) and 667MHz (1.6V). In parallel with Transmeta, AMD introduced their implementation of the variable voltage processor with *PowerNow!*, the Athlon 4 [14]. The latest version can run at five different voltage (1.2V – 1.4V) and frequency settings ($\leq 0.6\text{GHz}$ – 1.1GHz). Intel has introduced the *XScale* processor in 2001 [13]. This processor type is based on a StrongARM core and offers 16 different operational frequencies. According to the application specific standard product, the possible voltages/frequency settings vary. Only recently, Intel released the PXA800F processor [15] particular suitable for cellular phones. Based on an ARM core with *XScale*, this processor has the ability run at 104MHz and 312MHz . The increasing availability of DVS processors adds credibility to the practicability of the DVS technique. Certainly, DVS is becoming an important energy reduction technique.

2.2.2 Dynamic Power Management

Unlike DVS, dynamic power management (DPM) is already around for quite a while. The main strategy of DPM is the shutdown of idle system components [25, 87]. An advantage of DPM is its generality [22], which allows its usage not only for digital circuitry, but also for other system components such as displays, hard drives, and analogue circuits. DPM approaches often differ in the employed shutdown policies (or power management policies). For example, in its most aggressive strategy components are switched off im-

mediately when they become idle. A second strategy is the timeout-based policy, which switches off components after a fixed idle interval. This policy is well known from the advanced power management (APM) widely used in today's notebook computers. Nevertheless, since the restart of a component involves a time and power overhead to restore its fully functional state, such greedy policies might not result in power savings or may even increase the dissipated power [23]. Therefore, a careful consideration of the applied policies is necessary to achieve the highest possible power savings [22–24, 87]. The problem with most real-life systems is the uncertainty with which future events occur. The quality of a power management policy depends therefore on the accuracy with which the future behaviour of the system can be predicted, in order to start-up currently inactive components or to shut down currently active components at the right moment.

2.2.3 Dynamic Power Management versus Dynamic Voltage Scaling

DVS and DPM are both useful techniques that help to reduce the energy dissipation of an embedded system. However, one valid question, which has not been discussed so far, is: "Why should one use a rather complex technique like DVS to slow down processing elements if it is possible to switch off components during idle intervals?" This question can be illustratively answered using a simple example. For clarity reasons, the timing overheads for voltage scaling and power management are neglected here. Consider the following situation. A processing element (processor) performs a certain task τ_x in $20ms$ and dissipates a power of $500mW$ when running with $33MHz$ at a nominal supply voltage of $3.3V$ and a threshold voltage of $0.8V$. To meet the performance requirements of the application, the task needs to be repeated every $30ms$. Thus, between each consecutive execution of the task there are $10ms$ of idleness during which the processor can be deactivated (DPM). Under these circumstances the execution of task τ_x results in an energy consumption of $500mW \cdot 20ms = 10mJ$. Figure 2.5 illustrates this situation.

Now consider the execution using voltage scaling. The repetition time of $30ms$ allow to prolong the execution time of task τ_x from $20ms$ to $30ms$. That is, instead of shutting down the processor during the $10ms$ of idleness, the processor's performance is reduced from $33MHz$ to $22MHz$ ($33MHz / (30ms / 20ms)$). The lower frequency allows the reduction of the supply voltage from $3.3V$ to $2.61V$, according to Figure 2.3 and the tolerable increase in circuit delay of $33MHz / 22MHz = 1.5$. An exact equation to derive this volt-

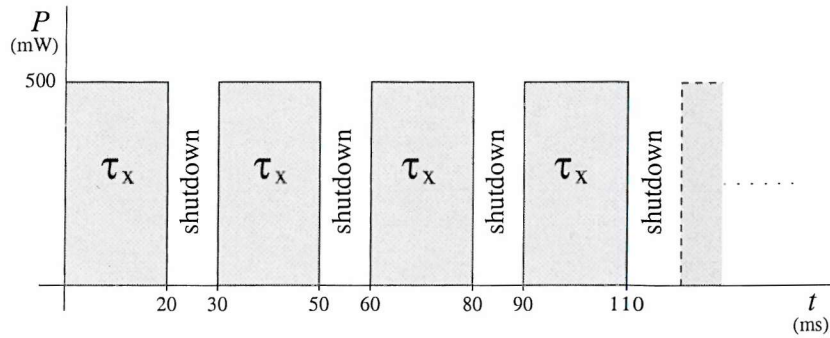


Figure 2.5: Shutdown during idle times (DPM)

age is given in Appendix D, Equation (D.5). At this voltage and frequency values the processing element dissipates a power of $209.8mW$ (see Figure 2.6). Thus, the energy

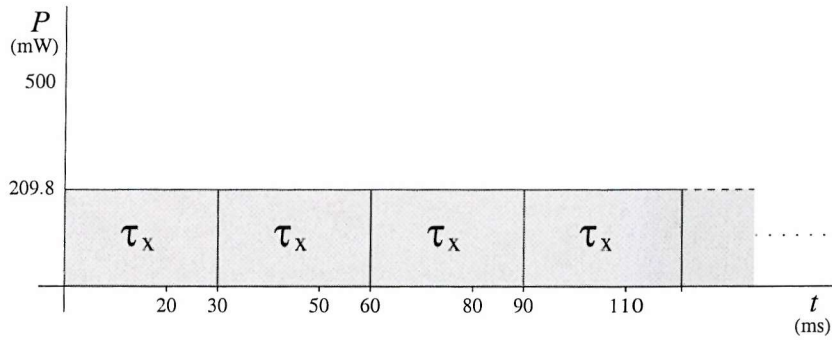


Figure 2.6: Voltage scaling to exploit the slack time (DVS)

consumption is given by $209.8mW \cdot 30ms = 6.29mJ$, a reduction of 32.1% compared to the $10mJ$ dissipated when using DPM. This simple example has shown that the energy efficiency of DVS is superior when compared to DPM. In fact, DVS always performs better than DPM, whenever both techniques are applicable [72].

2.2.4 DVS and DPM in Distributed Embedded Systems

Applying DVS and DPM to *distributed embedded systems* requires the careful consideration of interacting tasks. Consider the schedule shown in Figure 2.7, which was already introduced in Chapter 1 (Figure 1.9, Page 15). For instance, by slowing down CPU 2 during the execution of task τ_2 , the communications $\gamma_{2,4}$ and $\gamma_{3,5}$ are delayed, due to dependencies. This delay further influences the earliest possible start times of the tasks τ_4 and τ_5 . These interdependencies further complicate the voltage scaling processes [19, 64, 89, 122], compared to single processor systems. Another observation which can

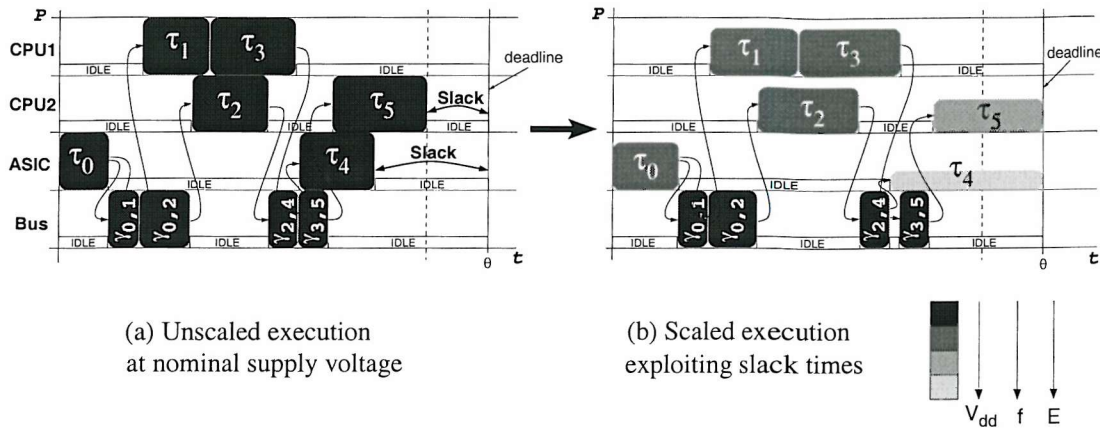


Figure 2.7: Combination of dynamic voltage scaling and dynamic power management

be made from Figure 2.7(b) is that even after applying DVS to exploit the available slack there remain some idle periods in the system schedule. Clearly, further lowering the voltage and frequency would result in missed deadlines θ , i.e., infeasible solutions. However, it is possible to switch off idle components during idle times, without influencing the schedule. In order to achieve a high degree of reduction, dynamic voltage scaling and dynamic power management should be considered together. However, due to the higher efficiency of DVS, DPM should be applied after DVS.

For DVS and DPM to be useful it is obligatory that the system experiences idle times (times where a certain components do not carry out an useful task) and slack times (times where a reduced system performance can be tolerated). Such idle and slack times can be found in distributed embedded systems due to four main reasons:

- It is often the case for a given application to show various degrees of parallelism, i.e., not all PEs will be utilised constantly during run-time. See, for instance, the schedule in Figure 2.7(a).
- The performance of the allocated architecture cannot be adapted perfectly to the application needs, since the allocation of "performance" is not given as a continuous range, but is rather quantised.
- The allocated architecture is commonly over-designed to allow an incremental design process, i.e., designers try to decrease development time through the reuse of the hardware architectures over several product generations. This is only possible by leaving enough "performance headroom" for future applications [107].

- (d) Schedules for hard real-time systems are constructed by considering worst-case execution times (WCETs), however, actual execution times of tasks during operation are, for most of their activations, smaller than their WCETs [131].

2.3 Energy Dissipation of Communication Links

The preceding sections mainly concentrated on the power consumption of processing elements and suitable energy management techniques. However, in distributed heterogeneous system with interacting components, the transfer of data between processing elements additionally contributes to the overall energy consumption. Fast communication is essential to avoid undesired contention of processing elements. Therefore, wide system buses (8 to 128 bit) and high-speed serial buses (CAN bus, I²C bus, USB, Gigabit Ethernet, Firewire, etc.) have become commonplace. With each transfer of data over the communication links (CLs), the line capacitance is charged and discharged, drawing a current from the I/O pins of the processing elements. The power dissipated by these currents is given by:

$$P_{CL} = \beta \cdot C_{bus} \cdot f_{bus} \cdot V_{tr}^2 \quad (2.12)$$

where $\beta \cdot C_{bus}$ is the effectively switched load capacitance of the bus lines, V_{tr}^2 is the operational voltage of the bus, and f_{bus} is the operational frequency at which bits are transferred over the bus. Important for the battery-lifetime, however, is the drawn energy:

$$E_{CL}^\gamma = N_{tr}^\gamma \cdot \beta \cdot C_{bus} \cdot V_{tr}^2 \quad (2.13)$$

where N_{tr}^γ denotes the number of bus cycles needed by communication γ . One possibility to reduce this energy dissipation is to encode the data before transferring it. Such bus-encoding techniques have been investigated with the aim of reducing the switching activity β [21, 102, 135]. Unlike the supply voltage of processing elements, the transmission voltage V_{tr}^2 cannot be reduced easily due to reliability issues, that is, environmental noise potentially corrupts data during transfers at low voltages. Nevertheless, many communication interfaces allow operation at different communication speeds (Ethernet at 10Mb/s, 100Mb/s, and 1Gb/s). In [86], Liu *et al.* present a technique that appropriately selects the speed of processors as well as communication links, in order to reduced the overall energy consumption. Furthermore, DPM can be equally applied to CLs as to PEs, i.e., during intervals where no data is transferred the CLs can be switched off.

2.4 Previous Work

This section briefly reviews the most relevant works in the area of co-synthesis for embedded systems, distinguishing between co-synthesis approaches that neglect issues related to power and more recent co-synthesis approaches that aim at the reduction of power consumption. Due to the topic of this thesis, special emphasis is placed upon the latter.

2.4.1 Co-Synthesis without Energy Minimisation

Hardware/software co-synthesis is a field of active research since the early 90's and several enlightening surveys have been published ever since [51, 93, 94, 136, 148]. Initially, most co-synthesis approaches targeted an architecture consisting of a single general-purpose processor connected to one or a few ASICs. Therefore, the main goal was to split the application between slow-but-cheap software and fast-but-expensive hardware, to achieve the highest possible performance without exceeding a given cost constraint [48, 49, 65, 80, 118, 145]. In these approaches special attention is given to different heuristic algorithms that solve the computational expensive partitioning problem, while issues related to power consumption are neglected. Ernst *et al.* [49] presented a simulated annealing based partitioning which starts from an all-in-software solution. Nodes are moved towards hardware until a given timing constraint is satisfied. Gupta and DeMicheli [65, 66] proposed a heuristic that iteratively moves nodes from hardware to software to reduce cost without the violation of imposed timing constraints. Kalavade and Lee [75] introduced a technique that utilises a global-criticality/local-phase measure to partition the system via a list scheduling approach [16]. Eles *et al.* [48] investigated a tabu-search partitioning algorithm, and a comparison to simulated annealing is provided. A genetic algorithm for hardware/software partitioning was developed in [118]. A comparison between simulated annealing, tabu-search, and genetic algorithm for one-CPU-one-ASIC architectures can be found in [145]. Other techniques that have been investigated include branch-and-bound algorithms [36] and dynamic programming [80].

More recent co-synthesis approaches target distributed heterogeneous architectures, i.e., architectures which potentially contain several processing elements and communication links of different types. The first work in this area was done by Prakash and Parker [108]. They formulated and solved the problem using mixed integer linear programming.

In a similar vein, Bender [20] applied mixed integer linear programming. In [149], Wolf developed a greedy heuristic which initially allocates a processing element for each task. Lightly loaded PEs are removed iteratively by re-assigning task to other components. Teich *et al.* [137] applied an evolutionary approach to the system-level co-synthesis problem using a graph-based mapping model. Oh and Ha [100] introduced a co-synthesis framework based on heterogeneous multiprocessor scheduling [99]. Pop *et al.* [106, 107] addressed co-synthesis problems that are typical to time-triggered communication protocols and incremental design processes. Madsen and Bjørn-Jørgensen [91] investigated the co-design of embedded systems under memory constraints.

2.4.2 Co-Synthesis with Energy Minimisation

All the approaches introduced in the previous section greatly neglect the optimisation of energy consumption. However, the recent development of the portable-application market has intensified the interest in system-level design techniques for energy-efficient embedded systems. The first approach that targeted the reduction of power dissipation throughout the co-synthesis process was proposed by Dave *et al.* [37]. They developed a constructive algorithm based on energy levels, which makes the mapping of tasks energy sensitive. Dick and Jha [45] reported a multi-objective genetic algorithm based co-synthesis approach. This framework simultaneously optimises the embedded system for cost, power consumption, and timing behaviour.

Most recently, co-synthesis approaches started to integrate the consideration of energy management techniques. This allows to optimise the embedded system towards the effective utilisation of DPM and DVS, hence, the reduction of energy consumption. In [70], Henkel introduces a low-power hardware/software partitioning approach for core-based systems. To avoid unnecessary high switching activity, the application is carefully partitioned into cores that can be selectively switched off. This technique is particularly suitable for designs in which no clock gating is applied. Lu *et al.* [88] presented a greedy on-line scheduling technique that optimises the execution order of tasks towards the utilisation of DPM. The main idea is to cluster the execution of tasks instead of scattering them, so that high shutdown overheads in terms of power and time can be avoided. Similarly, Brown *et al.* [30] developed a buffer insertion and scheduling technique for distributed systems that allows to reduce the shutdown overheads to improve the utili-

sation of DPM. In [72], Hong *et al.* introduced a design methodology for low power core-based real-time systems-on-a-chip. A variable voltage scheduling technique for single processor systems was proposed. Luo and Jha [89] developed a combined scheduling technique for periodic executing tasks with dependencies as well as aperiodic tasks. Dynamic voltage scaling is considered in this scheduling context. Dependent tasks are scheduled statically. The schedulability of aperiodic tasks as well as the utilisation of DVS are improved by distributing available slack time evenly among the processing elements of a distributed architecture. The on-line scheduler serves the aperiodic tasks and considers resource reclaiming. The same authors further enhanced their approach towards a battery-aware scheduling with the aim to improve the battery discharge profile [90]. Gruian and Kuchcinski [63, 64] extended a dynamic list scheduling heuristic to support DVS by making the priority function energy aware. In each scheduling step the energy-sensitive task priorities are re-calculated. If a scheduling attempt fails (exceeded hard deadline), the priority function is adjusted and the application is re-scheduled. Bambha *et al.* [19] presented a hybrid search strategy based on simulated heating, in order to derive an energy-efficient voltage selection for individual tasks. Mapping and scheduling are based on a dynamic list scheduling approach [132].

The techniques proposed in [19, 64, 89] have a close relationship to the problems that are addressed in this thesis. However, despite their energy-efficiency, these previous DVS approaches for heterogeneous distributed embedded systems have assumed a *fixed power dissipation* of processing elements during the energy minimisation. In reality this might not be true. For instance, low-level power minimisation techniques, such as clock gating [26, 27], result in power variations during the execution of an application. The problem of considering these power variations during dynamic voltage scaling is addressed in Chapter 3.

2.5 Concluding Remarks

This chapter has introduced the sources of power dissipation within distributed embedded systems. Two energy reduction techniques, namely dynamic power management and dynamic voltage scaling, have been outlined, and the advantage of DVS over DPM has been highlighted. Furthermore, an overview of most relevant previous works in the area of co-synthesis has been given.

Chapter 3

Power Variation-Driven Voltage Scaling

Dynamic voltage scaling is a powerful technique to reduce the energy consumption of processing elements within embedded systems. By simultaneously scaling supply voltage and operational frequency it becomes possible to trade off between performance and energy dissipation. This effect can be used to exploit the temporal performance requirement of an application, in order to save energy. Previous approaches to DVS in distributed systems have assumed that the power dissipation of processing elements is independent of the executed instructions [19, 64, 89]. Therefore, the voltage selection was carried out considering a constant power dissipation, in the following referred to as *fixed power model*. In practice this may not be true. For instance, modern IC designs often make use of low-level power minimisation techniques, such as clock gating to stop the switching activity in un-utilised blocks of the circuit [41, 140]; gated clocks are also used for DVS-PEs [31]. Under such circumstances the power dissipation can vary considerably during the execution of different functions [29]. This chapter presents a voltage scaling technique that takes this power variation effect into account, in order to increase the potential energy savings. The voltage scaling strategy is based on an energy-gradient driven heuristic, and the concept of mapped-and-scheduled task graphs is used to account for task and communication dependencies in a fast and effective way.

The remainder of this chapter is organised as follows. Section 3.1 motivates the need for a refined voltage selection that accounts for power variations. Section 3.2 introduces a new voltage scaling technique which accounts for power variations. Experimental results are presented in Section 3.3. Finally, concluding remarks are given in Section 3.4.

3.1 Motivation

This section motivates the consideration of power variations during the voltage scaling. The aim of this consideration is to improve the efficiency with which DVS can be applied.

Variations in the power dissipated by the individual tasks of an embedded system are caused due to the following two reasons: Firstly, modern IC designs make heavy use of gated clocks, i.e., unused circuitry is selectively "turned off" by stopping the clock signal to it [41, 140]; this holds also for DVS-PEs [31]. Thereby, unnecessary charging and discharging of the circuit load capacitance is avoided, which, in turn, results in lower power dissipation (Equation (2.7), page 25). For example, in the case of a general-purpose processor (GPP), including an integer unit and a floating point unit (FPU), it is not desirable to keep the FPU active if only integer instructions are executed. Hence, the clock signal to the FPU can be gated (stopped) during the execution of integer instructions, which will nullify the switching activity in the FPU. Thus, different tasks (different use of instructions) dissipate different amounts of power on the same processing element. In the case of an ARM7TDMI processor the supply current varies between $5.7mA$ and $18.3mA$, depending on the functionality which is carried out [29]. The second reason for the power variation effect is typical for core-based designs, where several different cores reside together on a single chip. Consider an ASIC accommodating four different cores: a FIR filter, an IDCT algorithm, a DES encrypt/decrypt unit, CORDIC (coordinate rotation digital computing) algorithm. Clearly, these four cores vary considerably in complexity. For instance, logic synthesis results show that the FIR filter requires approximately three times less area than the DES crypto unit [11]. This heterogeneity results certainly in different power dissipations, accordingly to which cores are active at a given time.

Taking this power variations into account during the voltage scaling improves the overall energy efficiency, since the available slack time is distributed more fairly among the tasks. To illustrate the influence of power variation effects on the voltage selection, a motivational example is given next. However, before starting with the example, it is necessary to define the term *energy-gradient*, which will be used throughout this chapter.

Definition 3.1 An *energy-gradient* ΔE_τ is defined as the difference between the energy dissipation of task τ with the execution time t_{exe} and the reduced energy dissipation (due to voltage and clock scaling) of the same task when extended by a time quantum Δt .

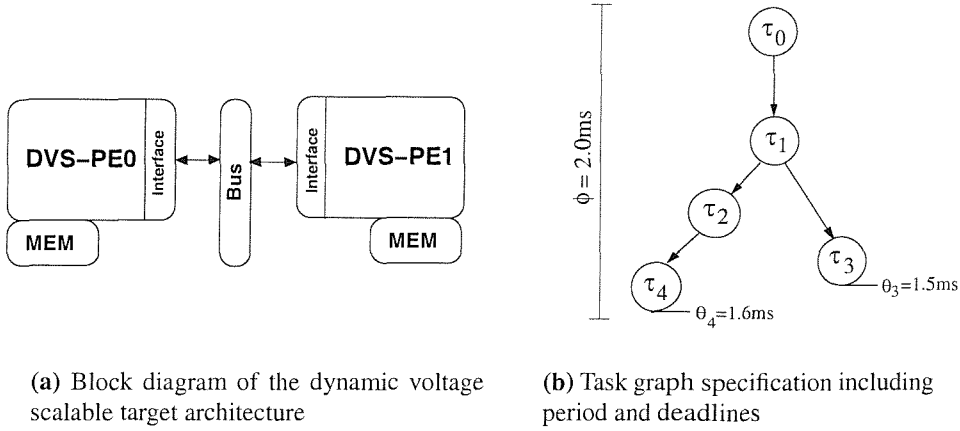


Figure 3.1: Architecture and specification for the motivational example

Mathematically:

$$\Delta E_{\tau} = E_{\tau}(t_{exe}) - E_{\tau}(t_{exe} + \Delta t) \quad (3.1)$$

where $E_{\tau}(t_{exe})$ and $E_{\tau}(t_{exe} + \Delta t)$ are calculated based on Equation (2.10). \square

Motivational Example: Considering Power Variations during Voltage Scaling

The intention with this illustrative example is to motivate the consideration of power variation effects during the voltage scaling of heterogeneous distributed systems. This is done by using two different power models during the DVS optimisation: (a) the traditional *fixed power model* (assuming constant power) [19, 64, 89] and (b) a *power variation model* which accounts for the power dissipation of each task (as proposed in this chapter). Additionally, in order to address the voltage scaling problem in the presents of power variations a new energy-gradient voltage scaling strategy is introduced.

The starting point for applying any DVS technique is a scheduled (at nominal voltage and frequency) system specification consisting of tasks and communication, which are mapped onto an allocated architecture. In this simple example, the considered architecture is composed of two heterogeneous DVS-PEs (e.g. a Transmeta Crusoe [78] and a StrongARM with Xscale technology [13]), as shown in Figure 3.1(a). Each PE has its own local memory to store the different tasks that are mapped onto it. These processing elements are connected through a single bus. The system specification is given by the task graph shown in Figure 3.1(b), including repetition rate ($\phi = 2ms$) and deadline constraints ($\theta_4 = 1.6ms$ and $\theta_3 = 1.5ms$). Nominal supply voltage V_{max} and threshold voltage V_t for

the two PEs are given in Table 3.1. This table further shows the nominal execution times

<i>task</i>	PE0 ($V_{max} = 5V$, $V_t = 1.2V$)		PE1 ($V_{max} = 3.3V$, $V_t = 0.8V$)	
	<i>exe. time (ms)</i>	<i>power dis. (mW)</i>	<i>exe. time (ms)</i>	<i>power dis. (mW)</i>
τ_0	0.15	85	0.70	30
τ_1	0.40	90	0.30	20
τ_2	0.10	75	0.75	15
τ_3	0.10	50	0.15	80
τ_4	0.15	100	0.20	60

Table 3.1: Nominal task execution times and power dissipations

and dynamic power dissipations of each task, according to their mapping (execution on PE0 or PE1). Furthermore, the transfer times and power dissipation of the communication activities are shown in Table 3.2, reflecting the inter-PE communications through the bus. Communications between tasks on the same PE (intra communications) are assumed to be instantaneous, and their power dissipation is neglected, as in most previous co-synthesis approaches [28, 37, 45, 80, 89, 105].

<i>comm.</i>	<i>comm. time (ms)</i>	<i>power dis. (mW)</i>
$\gamma_{0 \rightarrow 1}$	0.05	5
$\gamma_{1 \rightarrow 2}$	0.05	5
$\gamma_{1 \rightarrow 3}$	0.15	5
$\gamma_{2 \rightarrow 4}$	0.10	5

Table 3.2: Communication times and power dissipations of communication activities mapped to the bus

One possible mapping and scheduling of the system tasks and communications onto the underlying architecture is shown in Figure 3.2, which describes the power dissipation over time, that is, the power profile of PEs and CLs. It can be observed that PE0 accommodates task τ_0 and τ_4 , while the remaining tasks (τ_1, τ_2, τ_3) are mapped to PE1. In accordance to this mapping, two inter-PE communications over the bus are required between tasks τ_0 and τ_1 as well as between tasks τ_2 and τ_4 . The communication link CL0, connecting both PEs, shows these two communications, $\gamma_{0 \rightarrow 1} = (\tau_0, \tau_1)$ and $\gamma_{2 \rightarrow 4} = (\tau_2, \tau_4)$.

The dynamic system energy dissipation of this configuration at nominal supply voltage can be calculated using the dynamic power values and execution times given in Tables 3.1 and 3.2. The tasks mapped to PE0 (τ_0 and τ_4) consume an energy of $0.15ms \cdot$




Figure 3.2: Power profile of a possible mapping and schedule at nominal supply voltage (no DVS is applied)

$85mW + 0.15ms \cdot 100mW = 27.75\mu J$, while the tasks assigned to PE1 (τ_1, τ_2, τ_3) dissipate $0.3ms \cdot 20mW + 0.75ms \cdot 15mW + 0.15ms \cdot 80mW = 29.25\mu J$. Taking into account the communications on CL0 ($\gamma_{0 \rightarrow 1}$ and $\gamma_{2 \rightarrow 4}$), which consume $0.05ms \cdot 5mW + 0.10ms \cdot 5mW = 0.75\mu J$, the overall energy dissipation results in $27.75\mu J + 29.25\mu J + 0.75\mu J = 57.75\mu J$. Obviously, since the execution of task τ_3 finishes at $1.4ms$ and the task deadline is at $1.5ms$, a slack time of $0.1ms$ is available, as indicated in Figure 3.2. The same holds for task τ_4 , which finishes its execution after $1.5ms$, leaving a slack of $0.1ms$ until its deadline of $1.6ms$ is reached. These slacks can be used to extend the task execution times. Thus, the DVS-PEs can be slowed down by scaling the supply voltage and accordingly the clock frequency, following the relation given in Equation (2.8). Let us consider two cases for the identification of scaling voltages: (a) When a fixed power model is used (power variations are neglected, as in previous work [19, 64, 89]), i.e., all tasks mapped to the same PE are assumed to consume the same constant amount of power, and (b) a more generalised and realistic power model allowing for power variations among the tasks (as proposed in this work).

One approach to optimise the energy dissipation, which neglects the power profile, is to distribute the slack time evenly among the tasks, that is, in the given example each task is "stretched" using the same extension factor [64]. This is illustrated in Figure 3.3(a), where each task execution is extended using a factor of $e = 1.074$. For example, the

execution of task τ_0 was extended to $0.161ms$ ($0.15ms \cdot 1.074$). Similarly, the remaining tasks. The extension factor can be calculated considering the longest path to the task with the smallest slack time. In the example at hand, both deadline tasks (τ_3 and τ_4) have a slack of $0.1ms$; hence, we have the choice (both possibilities would result in the same extension factor). Consider the path indicated in Figure 3.2, which involves the tasks τ_0 , τ_1 , τ_2 , and τ_4 . The extension factor e is given by:

$$e = \left(\left(\sum_{\tau \in \mathcal{T}_p} t_{nom}(\tau) \right) + t_S \right) / \sum_{\tau \in \mathcal{T}_p} t_{nom}(\tau)$$

where t_{nom} is the nominal execution time, t_S denotes the slack time, and \mathcal{T}_p refers to all tasks on the path. Communications are neglected in this equation since they are not subject to scaling. Based on the execution values given in Table 3.1 and the $0.1ms$ slack time, the extension factor results in:

$$e = \frac{(0.15ms + 0.3ms + 0.75ms + 0.15ms) + 0.1ms}{0.15ms + 0.3ms + 0.75ms + 0.15ms} \approx 1.074$$

Thus, the extended task executions can be calculated as: $t_0 = 0.15ms \cdot e = 0.161ms$, $t_1 = 0.3ms \cdot e = 0.322ms$, $t_2 = 0.75ms \cdot e = 0.806ms$, $t_3 = 0.15ms \cdot e = 0.161ms$, and $t_4 = 0.15ms \cdot e = 0.161ms$. These execution times correspond to the schedule shown in Figure 3.3(a). In practice, extending a task is equivalent to executing the tasks at a lower performance, that is, at a lower operational frequency. This further allows to lower the supply voltages of PE0 and PE1 in accordance to the following formula, which has been derived from Equation (2.8) (see Appendix D for the derivation):

$$V_{dd} = V_t + \frac{V_0}{2d^*} + \sqrt{\left(V_t + \frac{V_0}{2d^*} \right)^2 - V_t^2} \quad (3.2)$$

where d^* denotes the normalised delay, which, in this example, is equal to the extension factor e . The constant V_0 is given by:

$$V_0 = \frac{(V_{max} - V_t)^2}{V_{max}} \quad (3.3)$$

Thus, the reduced voltages of PE0 is calculated as:

$$V_{dd} = 1.2V + \frac{(5V - 1.2V)^2 / 5V}{2 \cdot 1.074} + \sqrt{\left(1.2V + \frac{(5V - 1.2V)^2 / 5V}{2 \cdot 1.074}\right)^2 - (1.2V)^2}$$

$$V_{dd} = 4.788V$$

using the nominal supply voltage $V_{max} = 5V$ and the threshold voltage $V_t = 1.2V$ as given in Table 3.1. In the same way, the scaled supply voltage for PE1 can be calculated as $V_{dd} = 3.161V$, using $V_{max} = 3.3V$ and $V_t = 0.8V$. Adjusting the supply voltages of the PEs to these levels, the task deadlines are still satisfied, and the power dissipations are reduced. According to Equation (2.7), the power dissipation of each task can be calculated using the following relation:

$$\frac{P_{V_{dd}}}{P_{V_{max}}} = \frac{\alpha \cdot C_L \cdot f_{V_{dd}} \cdot V_{dd}^2}{\alpha \cdot C_L \cdot f_{V_{max}} \cdot V_{max}^2} = \frac{1}{e} \cdot \frac{V_{dd}^2}{V_{max}^2} = \frac{1}{d^*} \cdot \frac{V_{dd}^2}{V_{max}^2} \quad (3.4)$$

considering that $\alpha \cdot C_L$ is constant for a given task and that relation between reduced operational frequency $f_{V_{dd}}$ and maximal operational frequency $f_{V_{max}}$ is equivalent to the inverse of the extension factor $1/e = f_{V_{dd}}/f_{V_{max}}$. Out of this, the individual power dissipations can be calculated as: $P_{V_{dd}}(\tau_0) = 85mW \cdot 1/1.074 \cdot (4.788V)^2/(5V)^2 = 72.57mW$, $P_{V_{dd}}(\tau_1) = 20mW \cdot 1/1.074 \cdot (3.161V)^2/(3.3V)^2 = 17.08mW$, $P_{V_{dd}}(\tau_2) = 12.81mW$, $P_{V_{dd}}(\tau_3) = 68.33mW$, and $P_{V_{dd}}(\tau_4) = 85.38mW$. Taking into account the extended execution times and the energy dissipated by communications, the total energy consumption of the schedule shown in Figure 3.3(a) results in: $E = (72.57mW \cdot 0.161ms + 17.08mW \cdot 0.322ms + 12.81mW \cdot 0.806ms + 68.33mW \cdot 0.161ms + 85.38mW \cdot 0.161ms) + (5mW \cdot 0.05ms + 5mW \cdot 0.10ms) = 53.03\mu J$. This is equivalent to a reduction of 8.2%.

Now consider the case when the generalised power model (allowing power variations) is employed during the identification of scaling voltages for the task executions. This optimisation is based on an energy-gradient that has been defined in Equation (3.1). For a simpler illustration of the method, a fixed time quanta size $\Delta t = 0.01ms$ is assumed, which is 10 times smaller than the available deadline slack ($0.1ms$). Having defined the time quantum size Δt , it is now possible to calculate an energy-gradient ΔE_τ for each task, using Equations (3.2) and (3.4)). For instance, consider task τ_0 . The energy dissipation of this task at nominal supply voltage can be calculated as $E_0(0.15) = 0.15ms \cdot 85mW = 12.75\mu J$, using the values given in Table 3.1. Extending the execution

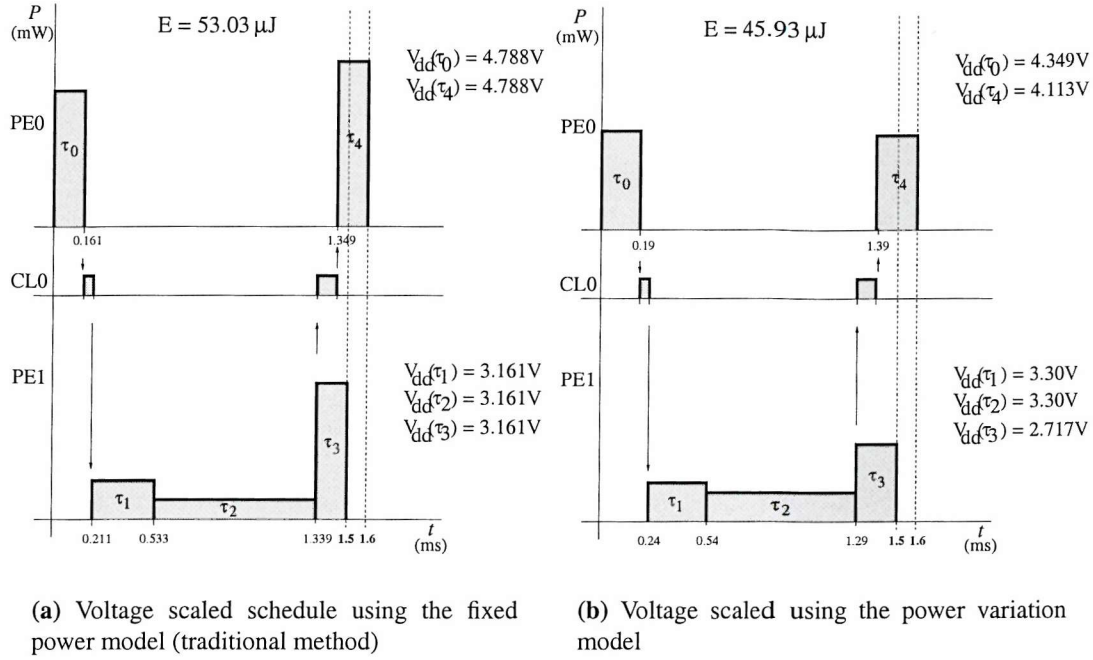


Figure 3.3: Two different voltage scaled schedules

of the task by $\Delta t = 0.01ms$ from $0.15ms$ to $0.16ms$ is equivalent to an extension factor of $0.16ms/0.15ms = 1.0667$. The scaled supply voltage of PE0 for this extension is $4.808V$ according to Equations (3.2). Running the task at this supply voltage would result in a power dissipation of $85mW \cdot 1/1.0667 \cdot (4.808V)^2/(5V)^2 = 73.69mW$ (using Equation (3.4)), and hence lead to an energy consumption of $73.69mW \cdot 0.16ms = 11.79\mu J$. Thus, the energy-gradient is given by $\Delta E_0 = 12.75\mu J - 11.79\mu J = 0.960\mu J$. The energy-gradients of the remaining tasks are calculated in the same way and result in: $\Delta E_1 = 0.234\mu J$, $\Delta E_2 = 0.156\mu J$, $\Delta E_3 = 0.899\mu J$, and $\Delta E_4 = 1.130\mu J$. Clearly, the task with the highest energy-gradient (in this case task τ_4) will improve the energy dissipation by the highest amount when extended by Δt . Therefore, the first time quantum is assigned to task τ_4 . In this manner, time quanta are iteratively distributed among the tasks until deadlines prevent further scalings. This optimisation process is illustrated through Table 3.3. The table shows the evolving energy-gradients for each of the five tasks and highlights the task to which a time quantum $\Delta t = 0.01ms$ is assigned, that is, the task with the highest energy-gradient. The first row holds the energy-gradients that have been calculated just above. As mentioned, task τ_4 gains most out of an extension by $\Delta t = 0.01ms$. Hence, the time quantum is assigned to task τ_4 , reducing its energy consumption by $1.130\mu J$ from $15\mu J$ to $13.870\mu J$. Having extended this tasks, its energy-gradient needs to be recal-

iteration	Energy-gradient ΔE (μJ)				
	τ_0	τ_1	τ_2	τ_3	τ_4
1	0.960	0.234	0.156	0.899	1.130
2	0.960	0.234	0.156	0.899	0.965
3	0.960	0.234	0.156	0.899	0.833
4	0.820	0.234	0.156	0.899	0.833
5	0.820	0.234	0.156	0.768	0.833
6	0.820	0.234	0.156	0.768	0.725
7	0.708	0.234	0.156	0.768	0.725
8	0.708	0.234	0.156	0.663	0.725
9	0.708	0.234	0.156	0.663	0.636
10	0.616	0.234	0.156	0.663	0.636
11	0.616	0.234	0.156	0.578	0.636
12	0.616	0.234	0.156	0.578	0.562
13	0.541	0.234	0.156	0.578	0.562
14	0.541	0.234	0.156	0.507	0.562
15				0.507	
16				0.451	
extension	0.04	0	0	0.06	0.06

Table 3.3: Evolution of the energy-gradients during voltage scaling

culated as $\Delta E_4 = E_4(0.16) - E_4(0.17) = 13.870\mu J - 12.905\mu J = 0.965\mu J$. The updated energy-gradient is shown in the second row of Table 3.3 (iteration 2). No time quanta have been distributed to the tasks τ_0 , τ_1 , τ_2 , and τ_3 , thus, their energy-gradients remain unchanged. Considering the second iteration, task τ_4 still achieves the highest energy saving when extended by $\Delta t = 0.01ms$, with $\Delta E_4 = 0.965\mu J$. Therefore, task τ_4 is extended by a second Δt and its energy-gradient is updated to 0.833. In the third iteration task τ_0 gains most of an extension and thus is extended. This iterative extension of tasks is repeated until no further extensions are possible without violations of task deadlines. This can be seen, for example, in the 15th iteration. Here the previous task extensions result in a situation where the task τ_4 just finishes in time, i.e., before its deadline is exceeded. Any further delay of task τ_4 or of any task that influences the finishing time of task τ_4 (i.e. the τ_0 , τ_1 , τ_2) would render the schedule infeasible. For this reason the tasks τ_0 , τ_1 , τ_2 , and τ_4 are removed from the list (Table 3.3). Similarly, the last remaining task τ_3 reaches its deadline after the 16th iteration. At this point, no further scaling is possible and the distribution of slack is aborted. Note, although time quanta of $0.01ms$ are distributed 16 times, only the available slack of $0.1ms$ is exploited. This is due to the fact that some task extensions require slack from both PEs (because of task dependencies), while others are

restricted to the slack of one PE. The last row of Table 3.3 shows how much the individual tasks have been extended. Accordingly, the new execution times are given as follows: $t_0 = 0.15ms + 0.04ms = 0.19ms$, $t_1 = 0.3ms + 0ms = 0.3ms$, $t_2 = 0.75ms + 0ms = 0.75ms$, $t_3 = 0.15ms + 0.06ms = 0.21ms$, and $t_4 = 0.15ms + 0.06ms = 0.21ms$. These extended execution times allow to lower the supply voltages during the execution of tasks τ_0 , τ_3 , and τ_4 to 4.349V, 2.717V, and 4.113V (using Equation 3.2), respectively, while the tasks τ_1 and τ_2 are ran at nominal supply voltage. Thus, the power dissipations of the tasks are: $P_{Vdd}(\tau_0) = 50.77mW$, $P_{Vdd}(\tau_1) = 20mW$, $P_{Vdd}(\tau_2) = 15mW$, $P_{Vdd}(\tau_3) = 38.74mW$, and $P_{Vdd}(\tau_4) = 48.33mW$. This results in a task energy dissipation of $E = 50.77mW \cdot 0.19ms + 20mW \cdot 0.30ms + 15mW \cdot 0.75ms + 38.74mW \cdot 0.21ms + 48.33mW \cdot 0.21ms = 45.18\mu J$. Adding the communication energy, the total energy consumption is given by $45.18\mu J + (5mW \cdot 0.05ms + 5mW \cdot 0.10ms) = 45.93\mu J$. This represents an energy reduction of 20.47%, which substantially higher compared to a reduction of 8.2% obtained with the power profile neglecting approach. Concluding, the power variations should be taken into account during the voltage selection in order to achieve further reductions in the energy dissipation. Previous DVS approaches for distributed system [19, 64, 89] have considered the fixed power model, i.e., power variations were neglected. To overcome this limitation, the following section introduces a new voltage scaling technique capable of taking power variations into account with the aim to improve energy savings.

3.2 Novel Algorithms for Dynamic Voltage Scaling

Having briefly outlined a new energy-gradient-based voltage scaling strategy in the previous section and shown the advantage of taking the power variations into account, this section introduces effective algorithms to implement this approach. Section 3.2.1 describes an novel energy-gradient voltage scaling, which is based on a new mapped-and-scheduled task graph structure (MSTG). In Section 3.2.2, this algorithm is enhanced towards discrete voltage selection. Section 3.2.3 analyses the computational complexity of the proposed voltage scaling algorithm.

3.2.1 Energy-Gradient-Based Voltage Scaling

The aim of the introduced DVS approach for distributed systems is to identify scaling voltages under the consideration of power variation effects. This is done for the scheduled and mapped system specification, such that the total dynamic energy dissipation is minimised. The presented approach assumes that no restrictions are placed on the scaling voltages, i.e., the technique targets variable-voltage systems (nearly continuous range of possible supply voltages) rather than multi-voltage systems (small and limited number of potential supply voltages). One particular aim of the voltage scaling technique is the energy estimation during the co-synthesis. Therefore, low optimisation times are of crucial importance, since the voltage scaling is executed in the innermost loop of an iterative co-synthesis process. One step towards this goal is the mapped-and-scheduled task graph structure, which will be introduced as part of the following voltage scaling algorithm. The overall voltage scaling algorithm is summarised in Figure 3.4, which is outlined next.

The input to the algorithm consists of the task graph specification, which has been mapped and scheduled beforehand onto a given system architecture. Furthermore, execution times and power dissipations are part of the architectural information, which also includes other necessary component properties, like the nominal supply voltage V_{max} , the threshold voltages V_t , etc. The minimal extension time Δt_{min} denotes the minimal time quantum to be distributed in each step of the algorithm. It is defined in order to speed up the determination of the voltage selection by preventing insignificant small extensions, which would lead to trivial power reductions.

To allow a fast and correct extension of task executions, which might influence other tasks and communications of the system due to interdependencies, it is beneficial to capture the schedule and mapping information into the task graph (line 01 in Figure 3.4). This can be achieved by generating a mapped-and-scheduled task graph (MSTG) structure, which is a transformed copy of the initial task graph. Figure 3.5 illustrates this transformation. The transformation consists of two steps, also given in the pseudo code of the MSTG_TRANSFORM algorithm (Figure 3.6): Firstly, all communications (edges) that are mapped to communication links are replaced by pseudo communication nodes and appropriate edges, thereby preserving the specified functionality. Secondly, all nodes mapped to a certain PE or CL are traversed in chronological order of execution and linked by pseudo-edges, if an edge does not already exist. In this way, the schedule and map-

Algorithm: PV-DVS**Input:** - task graph $G_S(\mathcal{T}, C)$

- mapping
- schedule
- architectural information
- minimum extension time Δt_{min}

Output: - energy optimised voltages $V_{dd}(\tau)$
 - dissipated dynamic energy E

```

01: MSTG_TRANSFORM //Generate MSTG from  $G_S$ 
02:  $\mathbb{Q}_E \leftarrow \emptyset$ 
03: forall ( $\tau \in \mathcal{T}_d$ ) {  $\Delta t_d(\tau) := t_d(\tau) - (t_s(\tau) + t_{exe}(\tau))$  }
04: forall ( $\tau \in \mathcal{T}$ ) { calculate  $t_e$  }
05: forall ( $\tau \in \mathcal{T}$ ) { if  $t_e \geq \Delta t_{min}$  then  $\mathbb{Q}_E := \mathbb{Q}_E + \tau$  }
06:  $\Delta t = \frac{\min t_e}{|\mathbb{Q}_E|}$ , if  $\Delta t < \Delta t_{min}$  then  $\Delta t = \Delta t_{min}$ 
07: for all ( $\tau \in \mathbb{Q}_E$ ) { calculate  $\Delta E(\tau)$  }
08: reorder  $\mathbb{Q}_E$  in decreasing order of  $\Delta E$ 
09: while ( $\mathbb{Q}_E \neq \emptyset$ ) {
10:   select first task  $\tau_{\Delta Emax} \in \mathbb{Q}_E$ 
11:    $t_{exe}(\tau_{\Delta Emax}) := t_{exe}(\tau_{\Delta Emax}) + \Delta t$ 
12:   update  $E_{\tau_{\Delta Emax}}$ 
13:   forall ( $\tau \in \mathcal{T}$ ) { update  $t_s, t_E$  and  $t_e$  }
14:   forall ( $\tau \in \mathbb{Q}_E$ ) { if  $(t_e(\tau) < \Delta t_{min}) \vee (V_{dd}(\tau) \leq V_t(\tau))$ 
                        then  $\mathbb{Q}_E := \mathbb{Q}_E - \tau$  }
15:    $\Delta t = \frac{\min t_e}{|\mathbb{Q}_E|}$ , if  $\Delta t < \Delta t_{min}$  then  $\Delta t = \Delta t_{min}$ 
16:   forall ( $\tau \in \mathbb{Q}_E$ ) { update  $\Delta E(\tau)$  }
17:   reorder  $\mathbb{Q}_E$  in decreasing order of  $\Delta E$ 
18: }
19: delete MSTG
20: return  $E_\Sigma$ , and  $V_{dd}(\tau), \forall (\tau \in \mathcal{T})$ 

```

Figure 3.4: Pseudo code of the proposed heuristic (PV-DVS) algorithm

ping are inherited into the task graph, and the influence of a task extension can be easily propagated through the system schedule by traversing the MSTG in a breadth-first order. Consider, for instance, the extension of task τ_2 . The initial task graph (Figure 3.5(a)) does not reveal if the extension of task τ_2 has an influence on any other task or communication, except on task τ_4 . However, the scheduling and mapping shown in Figure 3.5(b) indicates that an extension of task τ_2 will influence the subsequent task τ_3 . The MSTG shown Figure 3.5(c) captures this dependency through an pseudo edge between the tasks τ_2 and τ_3 . The main advantage of this tactic lies within the linear time complexity of the breadth-




Figure 3.5: Capturing the mapping and schedule information into the task graph by using pseudo edges and communication task

Algorithm: *MSTG_TRANSFORM*

Input: - task graph $G_S(\mathcal{T}, \mathcal{C})$
 - scheduling and mapping information
Output: - mapped and scheduled task graph (MSTG)

```

01: forall ( $\gamma \in \mathcal{C}$ ) { //insert communication tasks
02:   if ( $\gamma$  is mapped to any link  $\mathcal{L}$ ) {
03:     insert new communication task  $\tau_C$ 
04:     insert edge  $\gamma_t$  from task source( $\gamma$ ) to task  $\tau_C$ 
05:     insert edge  $\gamma_f$  from task  $\tau_C$  to task sink( $\gamma$ )
06:     remove  $\gamma$ 
07:   }
08: }
09: forall ( $\kappa \in \mathcal{K}$ ) { //insert pseudo edges
10:   predecessor task  $\tau_P \leftarrow \text{NIL}$ 
11:   traverse all task  $\tau_\kappa$  mapped to  $\kappa$ 
   in chronological order {
12:     if ( $\tau_P \neq \text{NIL}$ )  $\wedge$  ( $\gamma_{P \rightarrow \kappa} = \text{NIL}$ )
13:       insert edge  $\gamma_{P \rightarrow \kappa}$  from  $\tau_P$  to  $\tau_\kappa$ 
14:     else
15:        $\tau_P = \tau_\kappa$ 
16:   }
17: }

```

Figure 3.6: Pseudo code of task graph to mapped-and-scheduled task graph transformation

first search algorithm [59], which is leveraged to update the start and end times of all influenced activities (tasks as well as communications), due to voltage scaling.

After the initial task graph has been transformed into a MSTG, the DVS continues with line 02 (Figure 3.4), where a priority queue \mathbb{Q}_E is initialised. In order to identify all extendable tasks, the algorithm calculates the available slack times Δt_d of each hard deadline task \mathcal{T}_d (line 03). This slack is given by the difference between task deadline $t_d(\tau)$ and the task finishing time $t_S(\tau) + t_{exe}(\tau)$, where t_S and t_{exe} denote the task start and execution time, respectively. The algorithm then calculates the available slack times t_e of all tasks, taking into account the interrelation between tasks and communications (line 04). For this purpose an inverse breadth-first search algorithm is used to visit all nodes of the MSTG to inherit the slack time of influenced tasks. If a task has several successors, the smallest slack is inherited to ensure that no deadlines will be exceeded during the voltage scaling. In line 05, the algorithm includes tasks with an available slack time t_e greater or equal than the minimal extension time Δt_{min} into the priority queue \mathbb{Q}_E . In this way, tasks with negligible small or no extension possibility are excluded from the scaling process. The initial extension time Δt is calculated by dividing the smallest slack time among the extendable task ($\min t_e$) by the number of extendable tasks $|\mathbb{Q}_E|$ (line 06). This extension time, however, should not be smaller than the minimal extension time Δt_{min} (the selection of an appropriate Δt_{min} will be discussed in Section 3.3). It is now possible to calculate the energy-gradients of all extendable tasks, that is, the tasks in the queue \mathbb{Q}_E , using the Equations (3.1), (3.2), and (3.4), as indicated in line 07. In line 08, the priority queue \mathbb{Q}_E is reordered in decreasing order of the energy-gradients.

Now the distribution of slack starts. The algorithm iterates the steps between lines 09 and 18 until no extendable tasks are left in the priority queue \mathbb{Q}_E . In each of these iterations the algorithm picks the first element from the priority queue, the task which leads to the highest energy reduction (line 10). This task is then extended by Δt and the energy dissipation value is updated (lines 11 and 12) according to Equations (2.8) and (2.10). In line 13, the extension is propagated through the MSTG, since successor tasks might have been affected by the extension in terms of start time t_S , end time $t_E = t_S + t_{exe}$, and available slack time t_e . This propagation is carried out using a breadth-first search to visit all successors of the extended task, i.e., tasks and communications that might have been affected by the extension. In the next step (line 14), inextensible tasks are removed from the extendable task queue \mathbb{Q}_E if their available slack t_e is smaller than the minimal ex-

tension time Δt_{min} , or their scaled supply voltage V_{dd} is smaller or equal to the threshold voltage V_t . Taking into account the tasks in the new extendable queue, the time quantum Δt is recalculated (line 15), and based on this value, the energy-gradients ΔE are updated (line 16). The priority queue \mathbb{Q}_E is reordered according to the new energy-gradients (line 17). At this point, the algorithm either invokes a new iteration or terminates, based on the state of the extendable task queue. On algorithm completion, the MSTG copy is deleted and the scaling voltages for each task execution as well as the total dynamic energy dissipations are returned (lines 19 and 20).

3.2.2 Discrete Voltages

The algorithm, as described above, produces scaling voltages under the assumption that variable-voltage PEs are available that support continuous voltage scaling. However, it is possible to adapt the generated scaling voltages towards multi-voltage PEs, which are able to run at a restricted number of predefined voltages, as for example the case for the processor presented by Burd *et al.* [31, 33]. This processor has the ability to run at 15 different operational frequency (5 to 80 MHz) and voltage settings.

It has been shown in [74] that the two discrete supply voltages V_{d1} and V_{d2} , $V_{d1} < V_{dd} < V_{d2}$, around the continuous selected voltage V_{dd} are the ones which minimise the energy dissipation, under the assumption that the time overhead for switching between different voltages can be neglected. Thus, the proposed approach can be used for voltage selection on multi-voltage PEs. Given a task τ with execution time t_{exe} at the continuous selected voltage V_{dd} , then, in order to achieve minimal energy consumption, the same task τ will execute on the multiple voltage PE for t_{d1} time units at the supply voltage V_{d1} and for t_{d2} time units at supply voltage V_{d2} , where

$$t_{exe} = t_{d1} + t_{d2} \quad (3.5)$$

$$t_{d1} = t_{exe} \cdot \frac{V_{d1} \cdot (V_{dd} - V_t)^2}{(V_{d1} - V_t)^2 \cdot V_{dd}} \cdot \frac{\frac{V_{dd}}{(V_{dd} - V_t)^2} - \frac{V_{d2}}{(V_{d2} - V_t)^2}}{\frac{V_{d1}}{(V_{d1} - V_t)^2} - \frac{V_{d2}}{(V_{d2} - V_t)^2}}. \quad (3.6)$$

3.2.3 Algorithm Complexity

As mentioned in Section 3.2.1, the voltage scaling algorithm is intended to be used within the innermost loop of the co-synthesis where scheduling and mapping are iteratively op-

timised. Therefore, a moderate computational complexity is desirable in order to allow a thorough exploration of the design space in reasonable amounts of time. The complexity of the proposed PV-DVS algorithm can be derived as follows. The initial task graph can be copied and transformed into a MSTG in linear time. The WHILE loop (line 09 in Figure 3.4) is executed in the worst case $n \cdot m$ times, where n is the number of nodes $|\mathcal{T}|$ in the graph, since all tasks might be extendable. However, depending on Δt_{min} and Δt , tasks might be extended more than once, and m , for the worst case, is the maximum number of such extensions. The inner part of the WHILE loop shows the following complexities: The propagation of extensions takes $n + e$ in the worst case ($e = |C|$ is the number of edges in the graph), since all nodes and edges might have to be visited by the breadth-first search (line 13). Removing inextensible tasks, again, might take n steps. At most n steps are needed to determine the new extension time Δt . And finally, updating the extendable queue takes n operations (the queue is implemented as Fibonacci heap [55]). All other calculations inside the WHILE loop are executed in constant time. Therefore, the time complexity of the proposed PV-DVS algorithm is given as $\mathcal{O}(n \cdot m(4n + e))$. It should be noted that the extendable task queue \mathbb{Q}_E is progressively reduced from length n to 0. The reduction is not uniform since it might occur that suddenly (at the same time) many tasks become inextensible and are excluded from the queue. This, additionally, indicates that the complexity derived above is valid for the worst case. Furthermore, to account for DVS-PEs which run at discrete voltages, the suitable supply voltage have to be derived from the continuous voltage. This can be done in linear time. In summary, the overall computational complexity (including MSTG generation, PV-DVS, and discrete voltage calculation) is quadratic in the number of tasks and communications.

3.3 Experimental Results: Energy-Gradient based Dynamic Voltage Scaling

To demonstrate the efficiency and the applicability of the proposed generalised DVS technique (considering the power variation model) in reducing the energy dissipation of heterogeneous distributed systems containing power-managed PEs, numerous experiments and comparisons with previously published approaches [19, 89] have been carried out. The DVS algorithm outlined in this chapter has been implemented using C++ on a Pentium-III/750MHz Linux PC with 128MB RAM. The used benchmark examples

consist of 7 task graphs taken from previously published literature [19, 73] and 28 task graphs¹ automatically generated using TGFF [44]. These benchmarks are used to cover a wide spectrum of application diversity. The complexity of these task graph examples varies between 8 to 100 tasks with 7 to 151 edges. The amount of processing elements (PEs) and communication links (CLs) in the technology libraries varies between 4 and 16. The benchmarks used in this chapter are grouped into three major sets:

- (a) Our TGFF generated task graphs (tgff1–tgff25, tgff4_t, tgff4_fixed), and tgff17_m are mapped onto heterogeneous architectures containing power-managed DVS-PEs and standard PEs without DVS. Thus, these examples show varying power characteristics and component properties. The benchmark examples tgff4_t and tgff4_fixed are identical to tgff4 with slight modifications; tgff4_t denotes a task graph alternative with a critical tight deadline, while tgff4_fixed uses only DVS-PEs with a fixed power dissipation. Similarly, the example tgff17_m is a slightly modified version of tgff17 with a different hardware architecture.
- (b) The examples of Hou and Wolf [73] are two hypothetical task graphs (Hou and Hou_clustered). The task graph Hou_clustered represents the same functionality as Hou but the task graph is collapsed from 20 to 8 tasks. Since their technology library does not contain any DVS-enabled PEs, the given PEs are extended towards DVS capability (with $V_t = 0.8V$ and $V_{max} = 3.3V$). These examples also show different power dissipations (power variations) among the tasks, unlike the first benchmarks set.
- (c) The applications used by Bambha *et al.* [19] consist of two differently implemented fast Fourier transformations (fft1 and fft3), a Karplus-Strong music synthesis algorithm (Karp10), a quadrature mirror filter bank (qmf4), and a measurement application (meas). These *real-life* benchmarks of moderate size (12–28 tasks) use architectures composed of 2 to 6 identical DVS-PEs, assuming constant power dissipation. Supply voltages are between 0.8 and 7 volts. The throughput constraints and initial average power consumptions are calculated at a reference voltage of 5 volts.

The following experimental results are split into two sections. The first set of experiments concentrates on the influence of the power variation effect (Section 3.3.1), while the sec-

¹ Available from <http://www.ecs.soton.ac.uk/~ms99r/benchmarks.html>

and demonstrate the importance of an appropriate selection of the minimal extension time (Section 3.3.2).

3.3.1 Power Variations

To demonstrate the influence of power variations on the efficiency of DVS, the PV-DVS algorithm (Section 3.2) is compared to a power neglecting approach, i.e., a DVS approach that makes use of the fixed power model. The power neglecting approach (in the following referred to as EVEN-DVS) is based on the intuitive idea to distribute available slack time *evenly* among the processing elements, somewhat similar to the slack distribution idea used in [89].

Results of Benchmarks (a) and (b):

Table 3.4 shows a comparison between the EVEN-DVS (fixed power model) and the PV-DVS approach (power variation model) using the tgff1-tgff25 benchmarks as well as the two examples from Hou *et al.* [73] (Hou and Hou_clustered). In order to judge the complexity of the individual benchmark examples, the table gives the number of nodes and edges in the task graphs. The comparison between the two DVS approaches is carried out with respect to the energy dissipation when no DVS is employed, i.e., when the tasks are executed at nominal supply voltage (highest energy consumption). Consider, for example, benchmark tgff17, which consist of 29 tasks and 56 communications between tasks. The unscaled execution (NO-DVS) of the application dissipates an energy of $23459\mu J$. Using an even distribution of slack time (EVEN-DVS) this power consumption can be reduce to $20396.41\mu J$, a reduction of 13.06%. However, considering the power variations by using the PV-DVS algorithm it is possible to further reduced the energy to $18334.01\mu J$, when compared to the nominal energy dissipation a reduction of 21.85%. For all examples shown in Table 3.4 it is assumed that the mapping and schedule have been pre-determined, using a static mapping and a schedule generated by a mobility-based list scheduling technique [150]. Thus, the energy reductions are solely achieved through voltage scaling. As expected, both the EVEN-DVS and the PV-DVS technique reduced the energy dissipation of the systems in all cases (Column 6 and 9). It can be observed that the proposed DVS heuristic (PV-DVS) was able to further reduce the energy dissipation of all examples, when compared to EVEN-DVS. Due to the particular implementation

<i>Example</i>	<i>No. of tasks / edges</i>	NO-DVS <i>(nominal)</i>	EVEN-DVS <i>(fixed power model)</i>		PV-DVS (proposed) <i>(power variation model)</i>	
		<i>Energy Dissip. (μJ)</i>	<i>Energy Dissip. (μJ)</i>	<i>Reduction (%)</i>	<i>Energy Dissip. (μJ)</i>	<i>Reduction (%)</i>
tgff1*	8 / 9	355	193.49	45.50	112.87	68.21
tgff2	26 / 43	743224	722412.15	2.80	683954.54	7.97
tgff3	40 / 77	554779	410653.67	25.98	267651.03	51.76
tgff4	20 / 33	431631	402904.08	6.66	375914.03	12.91
tgff4_t	20 / 33	431631	412854.36	4.35	397201.93	7.98
tgff4_fixed	20 / 33	176723	142986.91	19.09	124905.61	29.32
tgff5	40 / 77	4187382	3963647.60	5.34	3767450.25	10.03
tgff6	20 / 26	1419124	1401605.68	1.23	1396445.06	1.60
tgff7	20 / 27	2548751	2289878.34	10.16	1951579.52	23.43
tgff8	18 / 26	1913519	1774151.42	7.28	1668485.33	12.81
tgff9*	16 / 15	996590	974159.01	2.25	918048.34	7.88
tgff10	16 / 21	69352	51263.60	26.08	46483.97	32.97
tgff11	30 / 29	4349627	4293736.56	1.28	4263279.98	1.99
tgff12	36 / 50	2316431	2243710.55	3.14	2212111.25	4.50
tgff13	37 / 36	2912660	2425431.77	16.73	2333338.86	19.89
tgff14	24 / 33	15532	13546.62	12.78	12479.41	19.65
tgff15	40 / 63	62607	62078.93	0.84	60334.62	3.63
tgff16	31 / 56	3494478	2913341.14	16.63	2518711.99	27.92
tgff17	29 / 56	23459	20396.41	13.06	18334.01	21.85
tgff17_m	29 / 56	153120	134988.02	11.84	113889.21	25.62
tgff18	29 / 57	1277704	1196851	6.38	665969	47.88
tgff19	14 / 19	5939	4713.59	20.63	4395.37	25.99
tgff20*	19 / 25	77673	48334.30	37.77	40280.98	48.14
tgff21	70 / 99	3177705	3175497.22	0.07	2658534.22	16.34
tgff22	100 / 135	5821498	5036657.40	13.48	4445545.63	23.64
tgff23*	84 / 151	11567283	10791880.89	6.70	10133912.03	12.39
tgff24	80 / 112	5352217	5349024.86	0.06	5238478.58	2.13
tgff25	49 / 92	5735038	5648816.00	1.50	5502681.64	4.05
Hou*	20 / 29	12265	10337.05	15.72	7474.55	39.06
Hou_clustered*	8 / 7	14546	12661.78	12.95	10270.32	29.39

*Components used for these examples consists of DVS-PEs only

Table 3.4: Comparison of the presented PV-DVS optimisation with the fixed power model using EVEN-DVS approach

of the DVS algorithm, which distributes slack evenly among the PEs (EVEN-DVS), also slack is allocated on non-DVS-PEs. Therefore, the higher energy reductions of the proposed DVS algorithm are due to two facts. Firstly, the EVEN-DVS allocates slack time on non-DVS-PEs. These times, of course, cannot be exploited to lower the power consumption. Secondly, the proposed DVS technique considers the power variations during the voltage scaling. This leads to better energy reductions (see Motivational Example 1, Section 3.1). To distinguish between both effects, architectures that consists of DVS-PEs only have been indicated in Table 3.4. In these examples, the higher energy reduction is solely achieved by taking the power profile into account. The remaining examples achieve

Example	No. of Nodes/ Edges	NO-DVS	PV-DVS (power variation model)		
		Energy Dissip.	Energy Dissip.	CPU time (s)	Reduction (%)
fft1	28/32	29600	18154	0.12	38.67
fft3	28/32	48000	36772	0.13	23.39
karp10	21/20	59400	47737	0.12	19.63
meas	12/12	28300	25732	0.10	9.07
qmf4	14/21	16000	12740	0.11	20.38

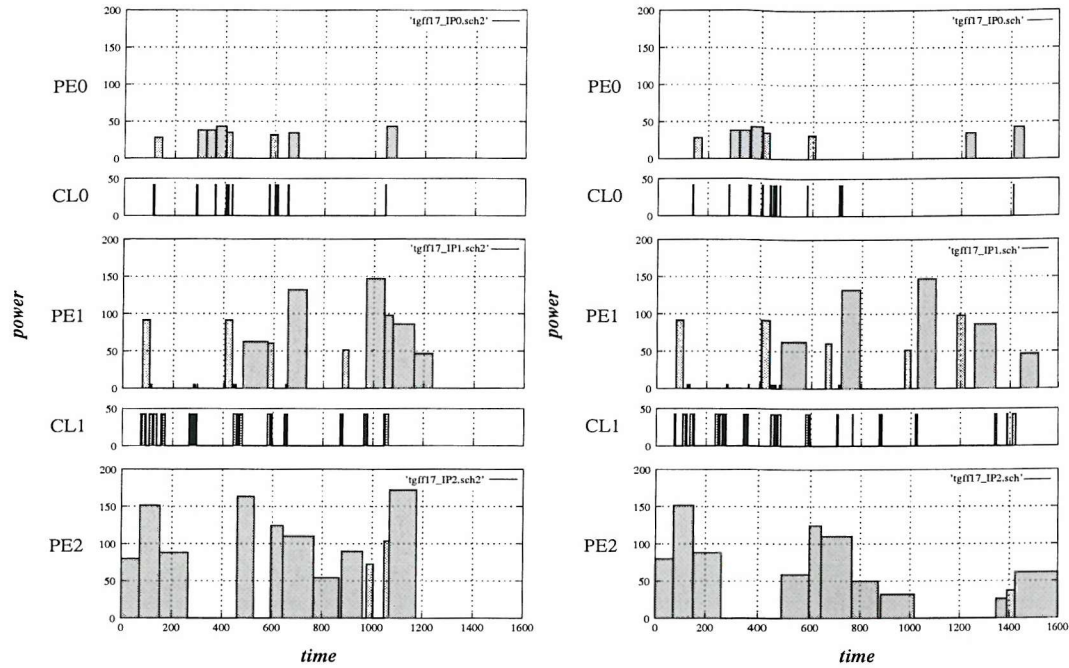
Table 3.5: PV-DVS results using the benchmarks of Bambha *et al.* [19]

increased energy efficiency due to both effects. It should be noted that it is hard to judge the achievable energy savings in a straightforward way, since the savings depend on task interdependencies as well as on the individual power dissipations of the given tasks.

In order to get an idea about the voltages scaled schedules consider Figure 3.7. This figure shows the task execution of benchmark `tgff17_m` for three different situations. Firstly, Figure 3.7(a) depicts the schedule at nominal supply voltage, i.e., tasks are executed as fast as possible ($f = f_{max}$ and $V_{dd} = V_{max}$), hence, consuming the maximal energy. Secondly, Figure 3.7(b) corresponds to an even distribution of slack time. And finally, Figure 3.7(c) illustrates the scaled execution based on the proposed PV-DVS algorithm. In this particular benchmark only one processing element (PE2) is DVS enabled, therefore, the execution properties of tasks mapped onto PE0 and PE1 are static.

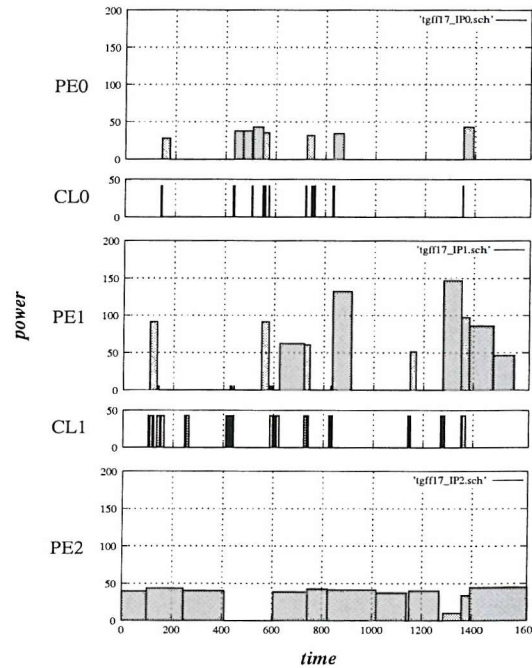
Results of Benchmark (c):

The next experiments are concerned with the benchmark set used by Bambha *et al.* [19]. Since they use a different communication model (contention, requests for the bus, etc.), the throughput constraints had to be re-calculated. Therefore, a direct comparison between the results reported in [19] and the results presented here is not valid. Nevertheless, the re-calculation of the throughput was carried out for the same task mapping and execution order as in [19], which is based on a dynamic level scheduling approach [132]. The results of these five examples, obtained using the PV-DVS method, are given in Table 3.5. It can be observed that in all cases the energy was decreased, with reductions between 9.07% to 38.67%. Furthermore, the highly serialised structure of `meas` allowed us to calculate the theoretically optimal voltage schedule for this example. Using the optimal supply voltages results in 13% energy saving. The PV-DVS algorithm achieved for the same example a reduction of 9.07%, which is only 3.93% higher than the theoretically



(a) Task execution at nominal supply voltage
(no DVS is exploited)

(b) Task execution at scaled supply voltages
using the EVEN-DVS approach (fixed power
model)



(c) Scaled task execution using the proposed
PV-DVS algorithm (power variation model)

Figure 3.7: Three identical execution orders of the `tgff17_m` benchmark: (a) unscaled execution at nominal supply voltage (NO-DVS), (b) using the EVEN-DVS, and (c) the PV-DVS approach




Figure 3.8: Energy reduction quality dependent on minimal extension time Δt_{min}

optimal solution. The reason for this for this is the greediness of the PV-DVS algorithm which always slacks the task which achieves the highest possible energy savings on a single PE; however, scaling this single task might require slack on several different PEs.

The presented results assume that computation and voltage scaling can be carried out concurrently, as the case of the processor introduced in [31]. Further, the time overhead needed by the processor to switch between two supply voltages is neglected since the used tasks are considered to be of coarse granularity. Therefore, the switching can be considered to be only a small fraction of the total task execution time (for real-world DVS processors this is in the range of 10–70 μ s for a full transition from the highest to the lowest supply voltage and vice versa [31]). However, in the case of fine-grained task this overhead might influence the optimisation and should therefore be considered.

3.3.2 Minimal Extension Time

To give insight into the dependencies between the computation effort, solution quality, and the minimal extension time Δt_{min} (see also the complexity analysis in Section 3.2.3), two additional experiments were conducted. In order to achieve accurate results, especially for the measurement of the optimisation times, the experiments are carried out using two large task graphs with 80 (tgff23) and 400 (large1) tasks. Figure 3.8 illus-




Figure 3.9: Execution time dependent on minimal extension time Δt_{min}

trates the dependency between the minimal extension time Δt_{min} and the solution quality (given as reduced energy E over nominal energy E_0). It can be observed that no energy reduction can be achieved until Δt_{min} is smaller than the largest slack available in the activity schedule (2364 for `tgff23` and 29581 for `large1`, see Figure 3.8). Clearly, if the algorithm must distribute time quanta bigger than any available slack, it cannot perform any voltage reduction, and therefore $E/E_0 = 1$. But energy reductions can be achieved by decreasing Δt_{min} to a value below the biggest slack that is present on a DVS-PE. In the case of example `tgff23`, the energy consumption approaches 87% of the nominal energy when reducing Δt_{min} below 2364. Nevertheless, it is not desirable to decrease the minimal extension time too much, since the additional reductions become insignificant (the curves level out) and the unnecessary extension will only increase the computational time of the optimisation.

The dependency between minimal extension time Δt_{min} and optimisation time of the DVS algorithm is illustrated in Figure 3.9, using a double logarithmic scale. It can be seen that with decreasing Δt_{min} the optimisation time increases. Similar to the observations given before, the algorithm cannot start any optimisation before Δt_{min} is smaller than the biggest slack given in the schedule (indicated by the vertical dashed lines in Figure 3.9). It is therefore important to find a good value for Δt_{min} , which trades off between solution quality and optimisation time. Finally, Figure 3.10 shows the energy reduction dependent




Figure 3.10: Energy reduction quality dependent on execution time

on the optimisation time. As the plot illustrates, with an increasing Δt_{min} the optimisation time decrease while the energy dissipation increases due to a less accurate calculation of the scaled supply voltages.

These experiment have shown that a well chosen Δt_{min} is essential for a fast and accurate energy estimation using the PV-DVS algorithm. Accordingly, a good choice is important to speed up the voltage scaling process. The following describes a heuristic method that has been used in the presented experiments to find an appropriate Δt_{min} setting for each solution candidate. It is based on the observation that for all conducted benchmarks experiments the characteristics shown in Figure 3.8 and Figure 3.9 hold. With reference to Figure 3.8, the nearly linear (in a semi-logarithmic scale) energy drop, after decreasing Δt_{min} below the highest DVS slack, is interpolated by a logarithmic function (indicated as "Interpolation" in Figure 3.8):

$$y = \alpha \cdot \log x + \beta \quad (3.7)$$

where the constants α and β are calculated using two initial points in the quasi linear part of the graph. The first point (Point 1) corresponds to the highest available slack s_h on any of the DVS-PEs, hence, matches the nominal energy dissipation. This point can be found in linear time and is indicated in Figure 3.8. To establish the second point needed for the interpolation, the DVS algorithm is run with a Δt_{min}^o three times smaller than the

highest DVS slack to find its corresponding reduced energy dissipation E° . For all used examples this was still in the steeply dropping part of the graph. The second point shown in Figure 3.8 was found in this way. Using both points, the constants α and β are given by:

$$\alpha = \frac{1 - E^\circ}{\log(s_h) - \log(\Delta t_{min}^\circ)} \quad (3.8)$$

and

$$\beta = E^\circ - \alpha \cdot \log(\Delta t_{min}^\circ) \quad (3.9)$$

The resulting linear interpolation for example `large1` is shown in Figure 3.8. Of course, finding the second point has a computational overhead, however, as it can be seen from Figure 3.8 and Figure 3.9, this "investment" pays off when compared to a wrong choice of Δt_{min} , which could result in much higher computational time or much higher energy consumption than necessary. The next step towards a good value for Δt_{min} is it to find a "rough" estimation for the achievable energy reduction E_e . For this purpose the average power dissipations on all DVS-PEs are calculated. The estimation used in the presented approach is based on the average power dissipation on each DVS-PE and the maximal available slack. Generally, such an estimation is too optimistic, i.e., it will be lower than the in reality achievable energy dissipation. An estimated energy dissipation for example `large1` is indicated in Figure 3.8. The minimal extension time Δt_{min} could be set to the intersection of the energy estimation and the interpolated energy drop. However, in the proposed heuristic the value of Δt_{min} is set one order of magnitude lower (as indicated by the cross in Figure 3.8). This is done to account for the fact that an energy estimation close to the real achievable energy reduction would be approximately one order of magnitude away from a good Δt_{min} . On the other hand, if the energy estimation would be far below the real achievable energy reduction, the calculated Δt_{min} would become unnecessary small. Therefore, no Δt_{min} smaller than 2.5 orders of magnitude (compared to the maximal DVS slack) is allowed. This is based on the observation that all used benchmarks show similar characteristics, and that an ideal Δt_{min} can mostly be found at maximal 2.5 orders of magnitude from the maximal DVS slack. It is fair to say, these extreme situations appear rarely during the optimisations process.

3.4 Concluding Remarks

This chapter has introduced a new energy minimisation technique for dynamic voltage scaling in distributed heterogeneous systems. The proposed technique is based on an energy-gradient based heuristic that takes into account variations in the power profile of DVS processing elements as well as the heterogeneity of the scalable components. In order to keep the computational complexity low, a novel mapped-and-scheduled task graph (MSTG) structure was proposed, which allows the propagation of task extensions in linear time. The selection of the minimal extension time was addressed through a heuristic approach that leverage observations taken from extensive experiments. Experimental results have reinforced the argument that power variations should be taken into account, in order to achieve higher energy saving compared to approaches that consider a fixed power model. Although the introduced PV-DVS algorithm has an overall quadratic computational complexity (while the fixed-power-model DVS can be applied in linear time), this complexity is low enough to scale system specifications of realistic size (10–100 tasks) in small amount of run-time.

Chapter 4

Optimisation of Mapping and Scheduling for Dynamic Voltage Scaling

Chapter 3 has shown how the energy consumption of distributed embedded systems can be minimised using a dynamic voltage scaling (DVS) technique that accounts for the power variations (PV) of tasks. This PV-DVS technique was applied to applications that had been statically mapped and scheduled beforehand. Nevertheless, the efficiency with which DVS can be applied depends significantly on the available idle and slack times within the system schedule. Clearly, application mapping (Section 1.3, page 10) as well as activity scheduling (Section 1.3, page 12) affect considerably these idle and slack times. For this reason a co-synthesis, which incorporates mapping and scheduling, should tightly integrate the consideration of DVS, in order to find solutions that carefully increase the amount of available idle and slack times on the DVS-PEs. This chapter introduces new techniques and algorithms for scheduling and mapping in distributed systems, which aim, in addition to traditional design goals such as performance and area usage, at an optimised utilisation of the DVS-PEs, and hence the reduction of the energy consumption. The overall co-synthesis flow is outlined in Figure 4.1. As indicated, the introduced co-synthesis split into two steps:

- Combined optimisation of scheduling and communication mapping
- Task mapping optimisation

The remainder of this chapter outlines this *two-step* co-synthesis process. The chapter is organised as follows. Section 4.1 concentrates on energy minimisation through sched-

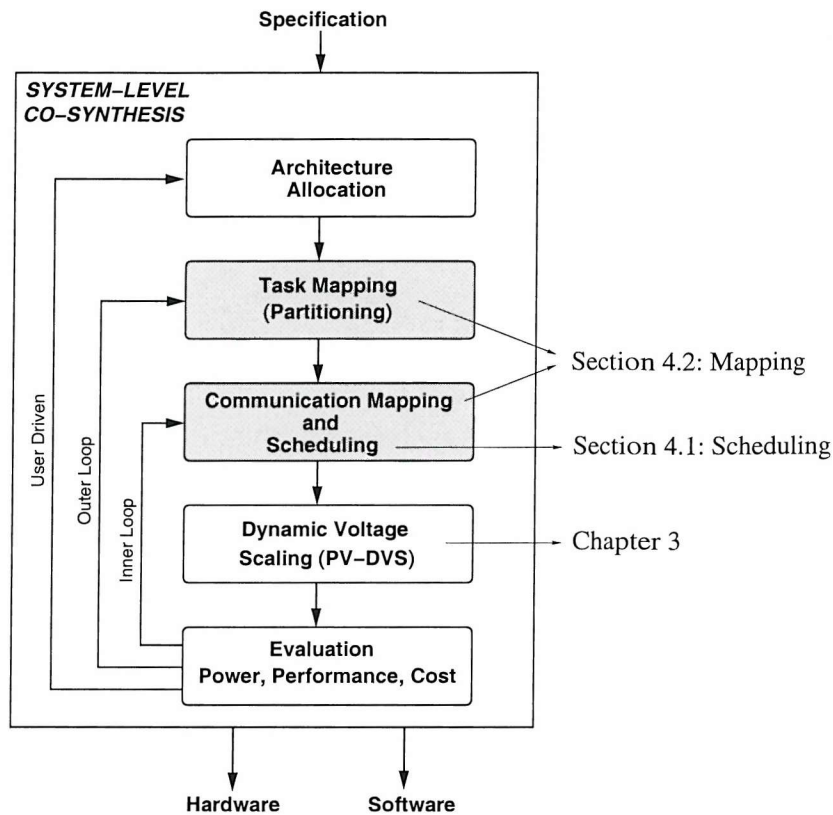


Figure 4.1: Co-synthesis flow for the optimisation of scheduling and mapping towards the utilisation of PV-DVS

ule optimisation towards an effective utilisation of PV-DVS. Techniques and algorithms for the optimisation of task and communication mapping are introduced in Section 4.2. Section 4.3 demonstrates through the usage of a real-life example how the proposed techniques can be applied in order to optimise the system implementation (including hardware architecture) towards energy-efficiency. Finally, Section 4.4 gives a summary of this chapter and draws some conclusions.

4.1 Schedule Optimisation

This section is concerned with the scheduling problem of heterogeneous distributed systems that contain DVS-PEs. The scheduling of tasks and communications greatly influences how efficiently DVS can be exploited, due to the direct impact on the available slack times. In general, the more slack is available in the schedule, the higher will be the achievable energy savings by exploiting DVS. This, however, becomes much more

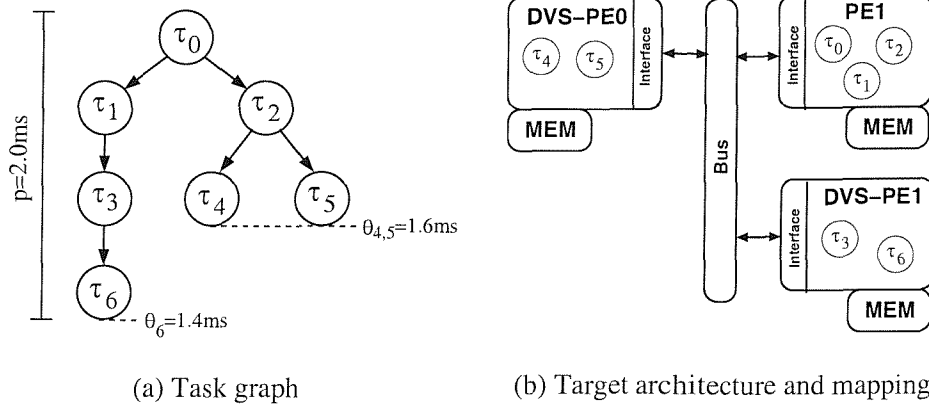


Figure 4.2: Specification and DVS-enabled architecture


complex and does not hold always for distributed systems under the consideration of the power variation model. In such a case, the available slack for high energy dissipating tasks should be considered more important than the slack of tasks consuming a minor amount of power.

This section first motivates the schedule optimisation under the power variation model in Section 4.1.1. This is followed by Section 4.1.2 which provides background information about scheduling techniques and genetic algorithms. Section 4.1.3 introduces a new scheduling algorithm for energy minimisation through PV-DVS. Finally, Section 4.1.4 presents experimental results.

4.1.1 Motivational Example: Scheduling

The purpose of this motivational example is to highlight the importance of taking power variations into account, while making energy-conscious scheduling decisions in the presence of DVS-PEs. Consider the system specification given as task graph in Figure 4.2(a). The seven tasks $\{\tau_0, \dots, \tau_6\}$ are mapped onto the architecture as shown in Figure 4.2(b). The architecture consists of two DVS-PEs (PE0, PE2) and one non-DVS-PE (PE1), which are connected through a single bus. The nominal supply voltage V_{max} and the threshold voltage V_t of PE0 and PE2 are $V_{max} = 3.3V$ and $V_t = 0.8V$, respectively, while PE1 runs all tasks at V_{max} (it cannot be scaled). The task execution times t_{nom} and power dissipations P_{max} at nominal supply voltage are given in Table 4.1, which also shows the task mapping. For the sake of simplicity, the communications are neglected when discussing this particular example. Figure 4.3(a) shows a possible schedule for the tasks executing at nominal

Task	exe. time t_{nom} (ms)	power dis. P_{max} (mW)	mapped to
τ_0	0.3	10	PE1
τ_1	0.3	20	PE1
τ_2	0.4	15	PE1
τ_3	0.1	40	PE2
τ_4	0.4	70	PE0
τ_5	0.2	90	PE0
τ_6	0.3	20	PE2

Table 4.1: Nominal execution times and power dissipations for the mapped tasks**Figure 4.3:** A possible schedule *not* optimised for DVS

supply voltage, i.e., at maximal supply voltage and consequently the highest energy dissipation. According to the values given in Table 4.1, the energy dissipation of this schedule can be calculated as $E = \sum_{\tau \in \mathcal{T}} (P_{max}(\tau) \cdot t_{nom}(\tau)) = 0.3ms \cdot 10mW + 0.3ms \cdot 20mW + 0.4ms \cdot 15mW + 0.1ms \cdot 40mW + 0.4ms \cdot 70mW + 0.2ms \cdot 90mW + 0.3ms \cdot 20mW = 71\mu J$. Considering the task deadlines given in Figure 4.2(a), it can be observed from the schedule in Figure 4.3(a) that the tasks τ_3 and τ_6 are eligible for scaling, since τ_6 finishes execution after $1ms$, while its deadline is at $\theta_6 = 1.4ms$. Hence, leaving a slack of $0.4ms$. An extension of any other task cannot be tolerated, since task τ_5 finishes execution just on its deadline $\theta_5 = 1.6ms$. By scaling the schedule, using the proposed implementation

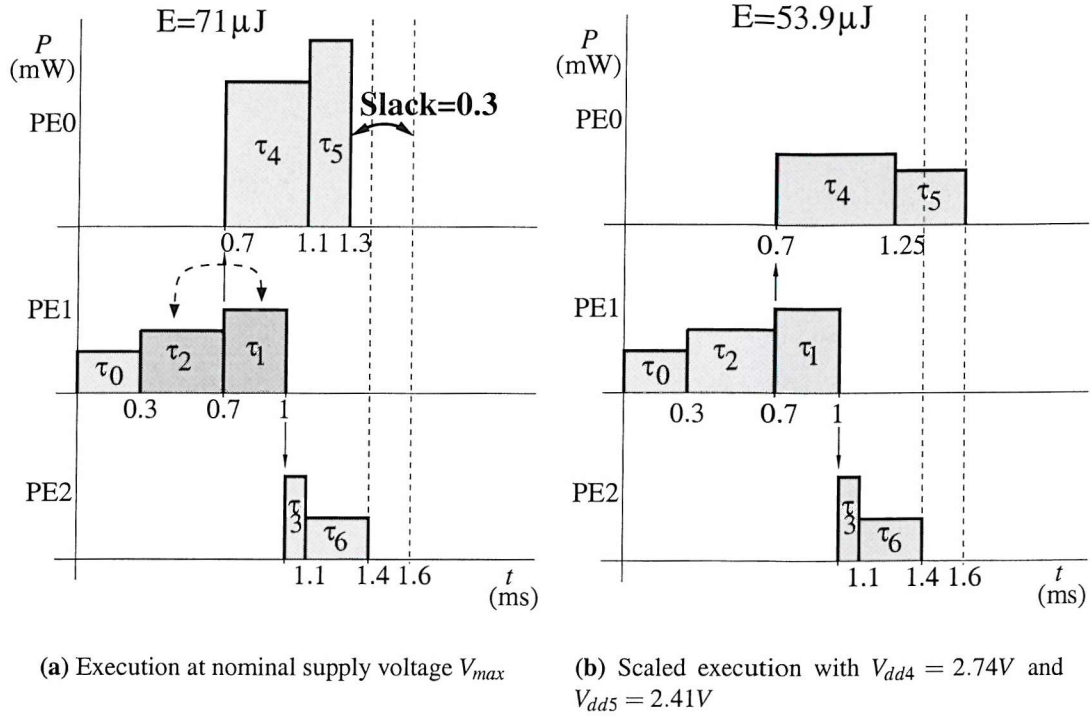


Figure 4.4: Schedule optimised for DVS considering the power variation model

of the PV-DVS technique (taking power variations into account) presented in Chapter 3, the voltage schedule shown in Figure 4.3(b) can be produced. In this scaled schedule the tasks τ_3 and τ_6 execute at 2.08V and 2.34V, respectively. Using Equation (2.6) and considering that the switched load capacitance $N_C^{\tau} \cdot \alpha \cdot C_L$ is constant for a given task, the energy consumptions of the tasks τ_3 and τ_6 are reduced to

$$E_{Vdd}^{\tau_3} = E_{max}^{\tau_3} \cdot \frac{N_C^{\tau_3} \cdot \alpha \cdot C_L \cdot V_{dd}^2}{N_C^{\tau_3} \cdot \alpha \cdot C_L \cdot V_{max}^2} = 0.1ms \cdot 40mW \cdot \frac{(2.08V)^2}{(3.3V)^2} = 1.59\mu J$$

and

$$E_{Vdd}^{\tau_6} = 0.3ms \cdot 20mW \cdot \frac{(2.34V)^2}{(3.3V)^2} = 3.01\mu J$$

Thereby, the total energy dissipation is given by $0.3ms \cdot 10mW + 0.3 \cdot 20mW + 0.4ms \cdot 15mW + 1.59\mu J + 0.4ms \cdot 70mW + 0.2ms \cdot 90mW + 3.01\mu J = 65.6\mu J$. This represents a reduction of 7.6%.

Now, consider a second feasible schedule at nominal supply voltage where the execution order of τ_1 and τ_2 has been swapped, as shown in Figure 4.4(a). This allows to start task τ_4 earlier, while the execution of task τ_3 is delayed. Since this re-scheduling does

not influence the nominal execution times and power dissipations, the nominal energy consumption remains $E = 71\mu J$ (as in first nominal schedule of Figure 4.3(a)). Observing the schedule reveals that task τ_6 just finishes execution within its deadline constraint $\theta_6 = 1.4ms$, while task τ_5 terminates execution after $1.3ms$. This leaves a slack of $0.3ms$ until its deadline $\theta_5 = 1.6ms$ is met. Therefore, only the tasks τ_4 and τ_5 can be scaled. Applying the PV-DVS technique of Chapter 3 returns the scaling voltages $V_{dd4} = 2.74V$ and $V_{dd5} = 2.41V$ for the tasks τ_4 and τ_5 , respectively. Executing the tasks at these voltages, the energy consumptions are reduced to $E_{V_{dd}}^{\tau_4} = 0.4ms \cdot 70mW \cdot (2.74V)^2 / (3.3V)^2 = 19.3\mu J$ and $E_{V_{dd}}^{\tau_5} = 0.2ms \cdot 90mW \cdot (2.41V)^2 / (3.3V)^2 = 9.6\mu J$. Hence, the total energy is reduced from $71\mu J$ to $0.3ms \cdot 10mW + 0.3 \cdot 20mW + 0.4ms \cdot 15mW + 0.1ms \cdot 40mW + 19.3\mu J + 9.6\mu J + 0.3ms \cdot 20mW = 53.9\mu J$. This is an improvement of 24.1%, while the first schedule achieved only a reduction of 7.6%. Note, although the schedule in Figure 4.4(a) shows less slack ($0.3ms$) than the one in Figure 4.3(a) ($0.4ms$), its energy reduction is significantly higher with 16.5%. Thus, more slack does not necessarily lead to increased energy savings. This is due to the particular power dissipations when executing the different tasks. In summary, this example has demonstrated how important it is to take into consideration the power variations during the schedule optimisation. In the presence of power variations it is not always true that an increased slack time results in lower energy dissipation, hence an energy-conscious scheduling technique should consider the power variation model.

4.1.2 Background

This section provides brief background information and introduces the terminology that is used throughout the rest of this chapter. This regards scheduling techniques (Section 4.1.2.1) as well as genetic algorithms (Section 4.1.2.2).

4.1.2.1 Scheduling Techniques

In general, scheduling techniques can be broadly classified into two categories: on-line (dynamic) scheduling techniques [40, 111] and off-line (static) scheduling techniques [16, 82, 99, 111]. The technique in the first class dynamically re-calculates the priorities of tasks during run-time of the application, i.e., the schedule can be changed during execution. Obviously, such approaches consume power and time during execution; time

which could otherwise be utilised by DVS to lower the energy dissipation. In the second class, a static schedule is calculated once before the application is executed (pre-run-time), i.e., the execution order of tasks and communications is maintained unchanged during run-time, hence, the power and time overhead is avoided. On-line scheduling are advantageous when no or only little knowledge is given about the tasks that have to be performed. In this case, the schedule can be dynamically adapted to the temporal needs of the application. However, embedded systems are often application specific, i.e., they execute a fixed set of tasks that is *a priori* known. Therefore, an adaption during run-time is not crucial for most embedded systems. In addition, off-line scheduling can guarantee that tasks deadlines are met during run-time, while this is not generally the case for dynamic scheduling [111]. Due to these reasons, the techniques introduced in this thesis concentrate on static scheduling.

Static scheduling for distributed systems that execute tasks with interdependencies has been intensively studied [16, 28, 82, 99, 108, 132, 150]. This scheduling problem belongs to the class of NP-hard problems [58], i.e., the search space for scheduling problems of realistic size is huge ($N!$, where N is the number of tasks and communication). Hence, finding an optimal solution is extremely computational expensive. For this reason most scheduling techniques rely on heuristic methods that produced not necessarily optimal solutions, but solutions of high quality. One of the most widely used scheduling heuristics is list scheduling (LS). List scheduling algorithms take scheduling decisions based on task priorities. They maintain one or more ready-lists which contain tasks that are ready to be scheduled. A static schedule is constructed by scheduling the ready task with the highest priority as soon as the eligible PE becomes available. Thereby, the assignment of priorities defines the task execution order. Most traditional LS approaches use various sophisticated algorithms to calculate these task priorities either statically [28, 150] (before list scheduling) or dynamically [81, 99, 132] (re-calculation after each scheduling step). Another method for the determination of priorities is used in genetic list scheduling algorithms (GLSA) [42, 62]. In contrast to list scheduling techniques, which produce a single schedule, GLSAs construct and evaluate many different schedules during an iterative optimisation process. However, most of this research on LS and GLSA concentrates on performances aspects only, i.e., the system schedule is solely optimised to achieve the highest possible throughput or to meet the imposed task deadlines. On the other hand, the genetic list scheduling approach presented in this chapter (Section 4.1.3) aims to find

schedules that simultaneously meet the system performance requirements and reduce the energy consumption by effectively exploiting the DVS-PEs. As the name suggests, genetic list scheduling approaches are based on genetic algorithms, hence, the following section briefly introduces the terminology and functionality of genetic algorithms.

4.1.2.2 Genetic Algorithms

Genetic algorithms (GAs) have been the subject of numerous investigations in the last decades, and they have been proven to solve different search and optimisation problems successfully [17, 60]. By imitating and applying the principles of natural selection and "survival of the fittest" on a population pool (several solution candidates), they are able to evolve (optimise) solutions to real-world problems. Each solution (individual) of the problem to be solved is encoded as a string (chromosome) and is associated with a solution quality (fitness). Based on their fitness, the individuals are ranked within the solution pool. In each iteration (generation) the algorithm selects upon the highest ranked individuals and gives them the opportunity to reproduce by mating (crossover) with different individuals of the population. This results in new individuals (offsprings) inheriting certain properties and features of the parent individuals, thus potentially increasing the probability to have higher solution quality. The produced offsprings replace the least ranked solutions in the population, which die out. New individuals are not only generated by means of crossover, but also by randomly changing (mutation) values (genes) of a chromosome, occasionally. This provides an additional opportunity to enter unexplored regions of the search space. The GA iterates until a certain stop criterion is fulfilled, for example, the maximum search time has been exceeded or no further improvement can be made.

4.1.3 Genetic List Scheduling Algorithm

In this section a new scheduling technique for energy minimisation through DVS is introduced, which addresses the problem motivated in Section 4.1.1, i.e., scheduling under the consideration of the power variation model. As mentioned in the beginning of this chapter, the presented co-synthesis approaches splits task mapping and scheduling into two separate optimisation steps, although some synthesis techniques combine these steps within a single optimisation [28, 42, 62, 75]. The decision to separate task mapping and

scheduling into two "independent" optimisations is based on the following two reasons:

- (a) The combination of list scheduling and mapping algorithms decide upon task priorities which task is to be scheduled next, but at this point it is *not* known where to execute the chosen task. Therefore, the execution time and power dissipation of the task are unknown as well. In this context, it is the duty of the scheduler to make a "greedy" mapping decision based on the power and time values with respect to the design objectives. However, DVS influences the execution times and power dissipations, hence, the mapping decision made upon the static values might prove to be wrong, especially from the energy reduction point of view. Separating the scheduling and mapping into two iterative optimisations overcomes this problem since the mapping is given before a schedule is constructed.
- (b) Due to the constructive nature of list scheduling and mapping algorithms a solution is constructed one by one. This results in a greedy approach, which is likely to get trapped at low quality or infeasible solutions in the presence of tight area and timing constraints. For instance, it is likely that the scheduler maps early tasks to fast hardware since hardware area is available. Nevertheless, the scheduling algorithm cannot look ahead and hence no area might be left when later, timing critical tasks arrive. A solution to overcome this problem was presented in [75]. However, this approach neglects issues related to power and a straight-forward enhancement towards DVS is not possible due to multiple competing design goals. Nevertheless, by splitting the problem into two steps, this greediness problem is avoided and the advantage of an increased search space, which is explored iteratively, can be leveraged. Clearly, increasing the search space results in higher optimisation times, however, it will be shown in Section 4.1.4 that these times are still reasonable.

The presented scheduling algorithm for the DVS problem under the consideration of the power variation model uses a genetic list scheduling approach to optimise the execution order of the tasks towards energy reduction and timing feasibility. It has been shown that the combination of genetic and list scheduling algorithms provides a powerful tool for the synthesis of multiprocessor systems. However, the proposed implementation varies in two fundamental issues from previous research [42, 62]:

- Instead of optimising the schedule solely for performance (reducing the make-

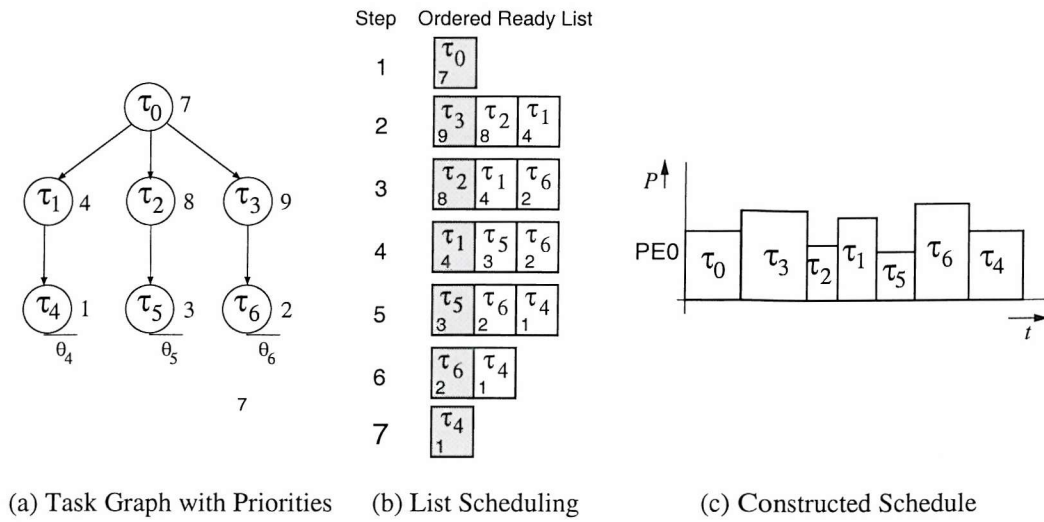


Figure 4.5: List scheduling

span), the proposed framework considers additionally the issue of energy minimisation with respect to DVS under the power variation model.

- The algorithms described in [42] and [62] employ a list scheduler which determines not only the execution order of tasks, but also their mapping. This combination is avoided in order to limit the greediness of the algorithm which would affect the solution quality.

As opposed to constructive list scheduling technique, genetic list scheduling approaches do not determine *one schedule* using a sophisticated algorithm for the priority assignment, but they construct and evaluate *many different schedules* during an iterative priority optimisation process. By encoding the task priorities into a priority string, it becomes possible to utilise genetic operators, such as crossover and mutation, to change task priorities and hence generate new scheduling solutions using static list scheduling. The principles behind genetic list scheduling approaches are outlined next.

Principle of Genetic List Scheduling

Genetic list scheduling approaches combine fast constructive list scheduling techniques with the optimisation power of genetic algorithms. The basic idea behind list scheduling is shown in Figure 4.5, which outlines the construction of a schedule for a single processor system. Consider the task graph with annotated priorities in Figure 4.5(a). In the initial




Figure 4.6: Task priority encoding into a priority string

scheduling step all tasks with no ingoing edges are placed into a ready list¹, as shown in Figure 4.5(b), Step 1. For the given task graph this is solely task τ_0 . Being the only task in the ready list, task τ_0 is scheduled. After finished its execution, the tasks τ_1 , τ_2 , and τ_3 become eligible for scheduling (due to their data dependency on τ_0); hence, they are placed into the ready list in decreasing order of their priorities (Scheduling Step 2). At this point τ_3 represents the task with the highest priority (9), hence it is scheduled in Step 2. Having scheduled task τ_3 , task τ_5 becomes ready and thus it is placed into the ready list. These scheduling procedure is repeated until no tasks are left in the ready list. Since each scheduling step schedules one task, seven iterations are necessary. The final schedule is shown in Figure 4.5(c). Clearly, different assignments of priorities result in different schedules. This is where the genetic algorithm comes into the play.

By encoding the task priorities into a priority string, as shown in Figure 4.6, it becomes possible to apply an genetic algorithm-based optimisation. The genetic algorithm aims to find an assignment of priorities that leads to a schedule solution of high quality in terms of timing behaviour and exploitable slack time. The main principles behind the genetic list scheduling algorithm are illustrated in Figure 4.7. These principles are based on two strategies: crossover and mutation.

Crossover Example

Out of an initial population pool that contains six priority candidate strings, the strings 1 and 3 are selected. Offsprings are produced by replacing parts of the first parent string with parts of the second parent sting. Hence, crossover results in two new offsprings (child 1 and child 2). These new priorities are used to schedule the activities, in order to

¹For distributed systems consisting of several processing elements and communication links, a ready list is introduced for each component. Furthermore, priorities are additionally assigned to communications.

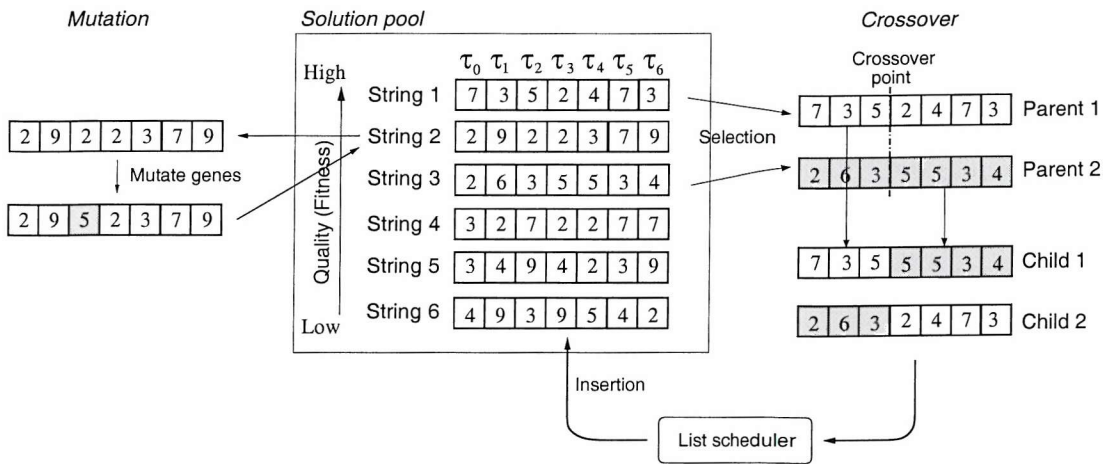


Figure 4.7: Principle behind the genetic list scheduling algorithm

determine their quality. According to the quality, the produced strings are inserted into the solution pool. By selecting high quality strings for crossover, the chances to evolve priority strings of even higher quality are increased. The mating of two strings is carried out with respect to an arbitrarily selected crossover point. □

Mutation Example

In order to enter into unexplored regions of the scheduling space, the genetic algorithm mutates individuals of the solution pool occasionally (with a low probability). The mutation is carried out by randomly changing genes of a randomly selected string. For instance, in Figure 4.7 string 2 is selected and its third gene is manipulated. The modified string is then reinsert into the solution pool. □

Both crossover and mutation are applied during the iterative execution of the genetic algorithm. The algorithm terminates after a stop criterion is fulfilled.

Genetic list scheduling approaches (GLSA) have several advantages compared to traditional (constructive) list scheduling methods. These advantages can be summarised as follows:

- Multi-objective optimisation is an important feature that is supported by genetic algorithms. It provides the opportunity to simultaneously optimise the implementation towards competing goals and so allows the system designer/architect to choose between several suitable implementations with different properties.
- The enlarged search space (at most $(|T| + |C|)!$ different schedules can be pro-

duced) provides the opportunity to find solutions of potentially higher quality.

- There is a large freedom to trade off between acceptable synthesis time and solution quality; as opposed to constructive techniques, where only one solution is produced rapidly. In addition, the search process can be initialised with such solutions that have been created by constructive techniques, and the GLSA can be used to further improve these.
- GAs with parallel populations and migration schemes provide a powerful approach to leverage additional computational power of computing clusters, which are becoming more and more widespread.

A drawback of any iterative improvement method for scheduling is the timing overhead involved in the successive construction of several scheduling solutions. Nevertheless, in timing critical scheduling situations the additional optimisation potential can be exploited to achieve timing feasible schedules and hence reduce the embedded system cost (this claim will be justified through extensive experiments in Section 4.2.3).

Implementation Details

The following introduces a DVS optimised genetic list scheduling algorithm (EE-GLSA²) that addresses the execution order of tasks and communications under the consideration of the power variation model. The pseudo code of the EE-GLSA is shown in Figure 4.8, which describes its features and implementation details. The solution pool (25 individuals) of the first generation is initialised (step 01) half and half by mobility-based [150] and randomly generated priorities (with values between the lowest and highest mobility), respectively. The mobility of a task is given by the difference between the as-late-as-possible (ALAP) start time and the as-soon-as-possible (start time) of a task [150]. This initial population was empirically found to be a good starting point, leading to fast convergence (i.e., low optimisation times). The algorithm then enters the main schedule optimisation loop (step 02–10), which is repeated until no improvement of at least 1% (with respect to the best found feasible schedule) is made within 10 generations. Each iteration of the loop goes successively through the following steps. All new priority candidate strings in the solution pool are used by the list scheduling algorithm to generate schedules at nominal

²Energy-Efficient Genetic List Scheduling Algorithm

Algorithm: <i>EE-GLSA</i>	
Input:	<ul style="list-style-type: none"> - task graph TG - mapping and execution properties corresponding to the mapping
Output:	- timing and energy optimised schedule
<p>01: Initialisation: Create initial population pool P of priority strings, half randomly generated and half based on mobility.</p> <p>02: Perform List Scheduling: Generates, for each member of the solution pool, a schedule based on the corresponding priority string.</p> <ul style="list-style-type: none"> a) Assign task priorities from the property string b) Invoke list scheduler without hole filling <p>03: Perform Voltage Scaling: Invoke the generalised DVS technique, calculating supply voltages for each task executed on a DVS-PE. This is done under the consideration of the individual power dissipation of tasks.</p> <p>04: Assign Fitness: Compute fitness of each individual in the population pool.</p> <ul style="list-style-type: none"> a) Calculate timing penalty b) Calculate energy based on the supply voltages c) Derive fitness based on energy and timing penalty <p>05: Termination: If no improved individual (improvement $> 1\%$) has been produced for 10 generations, then terminate. Otherwise, continue with step 06.</p> <p>06: Ranking: Individuals are ranked according to their fitness.</p> <p>07: Selection: According to the size of the generational overlap select individuals for mating. High ranked individuals have a high probability to be selected.</p> <p>08: Mating: Produce two-point crossover between a pair of selected individuals.</p> <p>09: Mutation: Randomly change genes of individuals using a dynamic mutation probability scheme, with exponential decreasing probability during run-time.</p> <p>10: Offspring insertion: Exchange low ranked individuals by newly produced individuals with respect to the size of the generational overlap. Continue with step 02.</p>	

Figure 4.8: Proposed EE-GLSA approach for energy-efficient schedules

supply voltage (step 02). The implemented list scheduler relies solely on the task priorities to make schedule decisions, i.e., no other techniques, like e.g. hole filling, are used to optimise the schedule. Although such techniques can improve the timing behaviour by eliminating idle periods in the schedule, the used list scheduler dissociate from them




Figure 4.9: Hole filling problem

since the DVS technique exploits exactly these idle times. For a simple example consider the task set given in Figure 4.9, mapped onto an architecture build out of two DVS-PEs. The task priorities are given on the right side of each task. According to these priorities, a static list scheduler can generate a feasible schedule, as shown in Figure 4.9. It can be observed that the tasks τ_0 , τ_1 , and τ_2 can be scaled only a small amount until deadline d_2 is met. However, task τ_4 can utilise the idleness before task τ_2 starts execution, and task τ_3 can be scaled until the deadline $d_{3,4}$ is met. Let us consider now the employment of a hole filling technique. In this case the list scheduler would try in its last scheduling step to place task τ_3 into the idle period between task τ_4 and task τ_2 . This decision, however, is fatally wrong from an energy reduction point of view, since the only available slack for all tasks would be the time between the execution of task τ_2 and the deadline d_2 , leaving not much headroom for voltage scaling. On the other hand, if the deadline $d_{3,4}$ would be identical with deadline d_2 , then the schedule displayed in Figure 4.9 would become infeasible and the schedule produced with hole filling would represent a better choice. To avoid such a dilemma, the used list scheduler solely makes scheduling decisions based on the task priorities, which are iteratively optimised. Therefore, the proposed EE-GLSA is capable of producing both schedule variants discussed above, and it evaluates their suitability with respect to timing and energy aspects to make an appropriate decision.

After the list scheduling has constructed a schedule (step 02), the algorithm proceeds by passing the built schedules to PV-DVS algorithm of Chapter 3, which identifies scaling voltages that minimise the energy dissipation (step 03). Note that schedules which exceed hard deadline constraints are still scaled as much as possible and are not excluded from the optimisation. This is done since good solutions are likely to be found as result of transformations performed on invalid configurations. However, a time violation penalty

is applied in such cases, as explained next. The scaled schedule is evaluated in terms of deadline violations and energy dissipation including the DVS reductions. Based on this evaluation, the fitness F_S of each schedule candidate is calculated (step 04) using the following equation:

$$F_S = \underbrace{\left(\sum_{\epsilon \in \mathcal{A}} E(\epsilon) \right)}_{\text{Energy diss.}} \cdot \underbrace{\left(1 + \frac{\sum_{\tau \in \mathcal{T}_d} DV_{\tau}^2}{T_{HP}^2} \right)}_{\text{Time penalty}} \quad (4.1)$$

where $\mathcal{A} = \mathcal{T} \cup \mathcal{C}$ defines the set of all activities (tasks as well as communications) and \mathcal{T}_d represents the set of all hard deadline tasks. The first part of the equation is used to calculate the total dynamic energy dissipation of all activities $\epsilon \in \mathcal{A}$. Based upon the type of activity, the energy dissipation can be calculated in the following way,

$$E(\epsilon) = \begin{cases} P_{max}(\epsilon) \cdot t_{nom}(\epsilon) \cdot \frac{V_{dd}^2(\epsilon)}{V_{max}^2(\epsilon)} & \text{if } \epsilon \in \mathcal{T}_{DVS} \\ P_{max}(\epsilon) \cdot t_{nom}(\epsilon) & \text{if } \epsilon \in \mathcal{T} \setminus \mathcal{T}_{DVS} \\ P_C(\epsilon) \cdot t_C(\epsilon) & \text{if } \epsilon \in \mathcal{C} \end{cases} \quad (4.2)$$

where P_{max} and t_{nom} refer to the power dissipation and execution time of tasks at nominal supply voltage, respectively. V_{dd} is the scaled supply voltage, \mathcal{T}_{DVS} represents the set of all tasks mapped to DVS-PEs, and P_C and t_C denote the power and execution time of communication activities. It should be noted that the power dissipations and the execution times of the tasks depend on the found scaling voltages V_{dd} , which have been calculated using the PV-DVS algorithm introduced in Chapter 3. The second part of the fitness function introduces a penalty factor due to deadline violations of deadline tasks,

$$DV_{\tau} = \max(0, (t_S(\tau) + t_{exe}(\tau)) - t_d(\tau)) \quad (4.3)$$

where $t_S(\tau)$ and $t_{exe}(\tau)$ denote the start time and the execution time (possibly scaled) of task τ , and $t_d(\tau)$ refers to the task deadline. T_{HP} represents the period of the hyper task graph (least common multiplier of all task graph periods), used to relate the deadline violation. Squaring has been applied in order to apply a higher penalty to larger violations of imposed deadlines.

By guiding the optimisation with this equation, the search for schedules is pushed into regions where low energy *and* feasible schedules are likely to be found. The algorithm then checks the halting criterion, as mentioned above. If the end of the optimisation has

not been reached, the algorithm continues (step 05) and the new priority candidates are ranked (step 06) and inserted into the solution pool based on their fitness values. Low ranked individuals of the pool are replaced (step 07) by new ones, which are generated through genetic crossover (step 08) and mutation (step 09). The employed GA is of steady state type, that is, not all of the individuals in the solution pool are replaced with each iteration (step 10). Steady state GAs were used due to their performance advantages compared to generational GAs, as indicated in [114]. The generational gap was set to 50%, i.e., half of the individuals in the solution pool survive unchanged in each generation. The crossover is carried out by means of a random two-point crossover [60]. To avoid a premature convergence towards suboptimal schedules, the idea of a dynamic mutation probability is leveraged [53]. This approach gives the algorithm the additional feature and capability to easily escape local minima in the beginning of the optimisation run. The mutation probability follows the equation $1/\exp(N_S \cdot 0.05)$ and is never allowed to drop below 15%. N_S denotes the current generation during the schedule optimisation. The generated offsprings are inserted into the population, resulting in a new generation. At this point, the next iteration is invoked and so different schedules are tried out.

4.1.4 Experimental Results: Schedule Optimisation

This section demonstrates through several benchmark experiments that the genetic list scheduling algorithm EE-GLSA introduced in this chapter achieves high energy savings, particularly when considering power variations during the optimisation. The EE-GLSA has been implemented on a PentiumIII/750MHz PC using a publicly available library of genetic algorithms [143]. All reported results represent average values that have been obtained over ten optimisation runs. In addition to the benchmarks (a–c) described in Chapter 3, the experiments conducted in this section also concern the benchmarks set used in [64].

- (d) Gruian's and Kuchcinski's task graphs [64] represent two sets (TG1 and TG2) of 30 randomly generated, communicating tasks with tight deadlines (determined by a critical path scheduling algorithm). These graphs show a high degree of parallelism and are mapped to architectures built of 3 or 10 identical DVS-PEs, assuming constant power consumption. These PEs are multi-voltage processor able to run at 3.3V, 2.5V, 1.7V, and 0.9V, while the threshold voltage V_t is 0.4V.

The following experiments are split into two sections. The first section concentrates on experiments that highlight the importance of considering power variations during scheduling in order to increase the energy-efficiency. While the second section compares the genetic algorithm based iterative schedule optimisation with constructive list scheduling techniques.

Power Variation Experiments

As demonstrated in the motivational scheduling example at the beginning of this section, power variations can have a significant influence on the suitability of a schedule when DVS is applied. This section assesses this influence through the usage of several benchmarks experiments. For this purpose Table 4.2 compares two different approaches; both are based on the same genetic list scheduling technique, however, the first approach considers the fixed power model, while the second employs the power variation model. The table shows for both approaches the achieved energy reductions and computational overheads when applied to the benchmark set *tgff*. The achieved energy reductions using the fixed power model varied between 1.35% in the case of *tgff6* and 51.89% in the case of *tgff3*. However, in all cases the energy could be further reduced when considering the power variations using PV-DVS. In these cases the reductions varied between 1.60% (*tgff6*) and 69.21% (*tgff3*). Accordingly, up to 17.32% higher savings could be achieved solely by reordering the execution of tasks and communications.

Clearly, an advantage of the EVEN-DVS technique (fixed power model) is its linear time complexity, when compared to the quadratic complexity of the PV-DVS algorithm (power variation model). This is also reflected in the reported optimisation time of Table 4.2. As it can be observed, the optimisation times for the approach using a fixed power model varied from 0.12s in the case of *tgff1* to 1.19s in the case of *tgff22*. The optimisation times for the same benchmarks optimised under the power variation model vary between 0.14s (*tgff1*) and 21.25s (*tgff22*). It is a classical trade-off between optimisation time and solution quality.

Comparison with Constructive List Scheduling Technique

In order to assess the quality of the new EE-GLSA technique to achieve high energy savings not only due to the consideration of the power variation model, this section provides

Example	GLSA + EVEN-DVS (fixed power model)			EE-GLSA (GLSA + PV-DVS) (proposed power variation model)		
	Energy (μ J)	CPU time (s)	Total Reduction (%)	Energy (μ J)	CPU time (s)	Total Reduction (%)
Tgff1	190.74	0.12	46.27	102.37	0.14	71.16
Tgff2	572920.30	0.20	22.91	545450.95	0.27	26.61
Tgff3	266906.88	0.33	51.89	170837.88	3.11	69.21
Tgff4	377444.70	0.24	12.55	375778.35	0.97	12.94
Tgff4_t	405473.20	0.25	6.06	396579.23	0.62	8.12
Tgff4_fixed	127867.31	0.26	27.65	124419.18	1.00	29.60
Tgff5	3721136.80	0.37	11.13	3450291.60	2.41	17.60
Tgff6	1399967.70	0.23	1.35	1396445.00	0.25	1.60
Tgff7	1924999.80	0.20	24.47	1797520.40	0.27	29.47
Tgff8	1722056.40	0.19	10.01	1648321.50	0.20	13.86
Tgff9	829607.95	0.18	16.76	774993.70	0.26	22.24
Tgff10	45324.61	0.17	34.65	44529.05	0.22	35.79
Tgff11	3755206.40	0.22	13.67	3621739.60	0.42	16.73
Tgff12	2212405.00	0.34	4.49	2198978.20	3.73	5.07
Tgff13	2342891.80	0.28	19.56	2315765.60	0.80	20.49
Tgff14	11891.05	0.21	23.44	11752.85	0.25	24.33
Tgff15	61271.38	0.41	2.13	60129.32	1.07	3.96
Tgff16	2492364.60	0.26	28.68	2449747.20	0.55	29.90
Tgff17	18922.85	0.27	19.34	18249.41	0.56	22.21
Tgff18	1724421.20	0.14	6.87	1421224.40	0.16	23.25
Tgff19	4514.75	0.17	23.98	4357.35	0.16	26.63
Tgff20	42703.58	0.19	45.02	37222.68	0.60	52.08
Tgff21	2983043.80	0.56	6.13	2578046.00	3.92	18.87
Tgff22	4664876.00	1.19	19.87	4119749.20	3.44	29.23
Tgff23	9826644.00	0.64	15.05	8855575.20	21.25	23.44
Tgff24	5240976.80	0.59	2.08	4881187.60	10.48	8.80
Tgff25	4922084.80	0.39	14.18	4545249.60	1.91	20.75

Table 4.2: Experimental results obtained using the *fixed power model* and the *power variation model* during voltage selection; both integrated into a genetic list scheduling algorithm

a comparison with constructive list scheduling techniques. The first comparison of the EE-GLSA is against a mobility-based (i.e., performance-driven) constructive list scheduling technique [150]. The results presented in Table 3.4 (page 55) have been generated using a mobility-based list scheduling, while the results Table 4.2 are based on the genetic list scheduling approach. Hence, comparing the results of both tables is equivalent to a comparison between both scheduling approaches. Take, for instance, example *tgff2*. When considering a constructive list scheduling based on mobility and a voltage scaling based on PV-DVS, an energy reduction of 7.97% can be achieved (Column 7 in Table 3.4). Nevertheless, applying the genetic list scheduling approach combined with PV-DVS, the energy is reduced by 26.61% (Column 7 in Table 4.2). Note, these savings are solely

achieved through an improved execution order of tasks and communications, since in both cases the same voltage scaling technique (PV-DVS) has been applied. Overall it can be observed that the genetic list scheduling technique yields higher energy savings in most of the examples. In fact, only in the case of benchmark `tgff6` no further improvement could be achieved through genetic list scheduling. This is due to the timing critical execution of tasks, which offers no energy reduction potential. Certainly, the GA based schedule optimisation introduces a computational overhead, which results in a necessary trade-off between solution quality in terms of energy dissipation and synthesis time. Nevertheless, the EE-GLSA can produce solution of at least the same quality as the mobility-based approaches in the "same" time, by placing a string with priorities based on the mobility criterion into the initial population of the EE-GLSA.

To further confirm the quality of the EE-GLSA, it is next compared to the DVS scheduling technique proposed by Gruian *et al.* [64]. This comparison is carried out using the benchmark collections TG1 and TG2, which contain 60 task graph examples with tight deadlines. The individual graphs capture highly parallel tasks with relatively few communications. Nevertheless, for experimental purpose these benchmarks represent a valuable choice. The reported energy reductions in [64] for these benchmarks are 28% and 13%. Table 4.3 presents the results obtained using the genetic list scheduling algorithm EE-GLSA, which is driven by the PV-DVS algorithm. The table separates both benchmark collections. Although the examples do not allow the EE-GLSA approach to leverage power variations, since the specified power values are constant, the achieved average energy reductions for TG1 and TG2 are 41.16% and 18.82% (Column 5 and 11, last row), respectively. This is an improvement of 13.16% and 5.82%, which indicates the effectiveness of the proposed optimisation technique, even when using constant power benchmark examples. However, since the results in [64] are obtained using multi-voltage PEs rather than variable-voltage PEs, additional experiments have been conducted, using the same discrete PE voltages ($V_{dd} = \{0.9V, 1.7V, 2.5V, 3.3V\}$, while $V_t = 0.4V$). As outlined in Section 3.2.2, each continuously selected voltage (using the PV-DVS algorithm) can be split into its two neighbouring discrete voltages considering the available voltages of the multi-voltage PE. The corresponding run-times at each voltage are calculated using the Equations (3.5) and (3.6). The results of the discrete voltage optimisation are shown in Table 4.3, see columns with the headings "Discrete Reduc.". For the two benchmark sets the achieved average energy reductions are 37.61% and 15.83%, respectively, which rep-

<i>TG1 example set</i> 10 Processors each executing 3 tasks						<i>TG2 example set</i> 3 Processors each executing 10 tasks					
Example	No. of nodes & edges	Contin. Energy Dissip.	CPU time (s)	Contin. Reduc. (%)	Discrete Reduc. (%)	Example	No. of nodes & edges	Contin. Energy Dissip.	CPU time (s)	Contin. Reduc. (%)	Discrete Reduc. (%)
r000	30/24	474783	6.09	40.56	37.35	r000	30/22	659972	0.41	20.30	17.24
r001	30/19	395964	10.97	47.87	44.11	r001	30/23	809191	0.55	13.99	11.39
r002	30/23	523822	2.97	29.67	26.71	r002	30/14	522578	1.76	25.94	24.35
r003	30/20	504252	12.67	49.31	45.83	r003	30/22	604127	1.06	20.97	18.11
r004	30/18	541220	4.32	38.98	35.13	r004	30/23	523668	1.18	25.27	22.23
r005	30/17	432787	8.27	41.89	39.10	r005	30/17	694278	0.94	20.40	17.13
r006	30/16	592832	4.33	34.25	31.03	r006	30/16	632867	1.64	20.27	17.66
r007	30/23	551119	6.30	34.23	31.07	r007	30/23	717256	0.42	7.94	5.34
r008	30/20	590249	4.20	31.56	28.69	r008	30/18	513749	1.70	34.06	30.32
r009	30/14	369454	6.98	45.76	41.71	r009	30/18	763976	0.81	14.80	10.92
r010	30/25	315060	13.43	55.35	51.48	r010	30/18	669359	0.91	21.04	17.66
r011	30/23	421029	11.36	47.61	43.83	r011	30/21	644326	0.90	20.79	16.58
r012	30/20	675304	5.51	32.11	29.06	r012	30/18	660552	0.29	13.03	9.80
r013	30/20	543053	5.87	33.64	29.87	r013	30/20	697671	0.34	10.45	7.64
r014	30/19	429695	12.10	50.73	46.36	r014	30/22	658146	0.35	17.09	14.12
r015	30/17	420428	11.15	48.31	44.53	r015	30/15	738860	0.85	13.34	9.85
r016	30/24	481347	5.99	35.37	32.21	r016	30/20	538039	1.32	22.12	18.98
r017	30/27	510259	13.37	42.47	38.97	r017	30/16	654964	1.53	22.29	19.42
r018	30/17	426713	9.70	39.10	36.08	r018	30/20	706124	0.90	19.04	16.91
r019	30/23	577211	7.18	33.82	29.61	r019	30/19	693263	1.56	19.15	15.80
r020	30/24	395536	10.41	46.89	42.66	r020	30/23	684600	0.58	16.34	13.39
r021	30/15	675748	5.49	31.39	28.64	r021	30/20	667156	0.29	13.28	10.44
r022	30/24	457843	13.59	44.05	40.63	r022	30/22	454777	1.15	33.71	29.94
r023	30/20	379382	11.78	52.50	47.87	r023	30/18	579582	0.96	18.43	15.59
r024	30/16	530941	6.23	35.88	31.70	r024	30/19	633231	0.86	22.15	18.89
r025	30/15	301150	15.67	60.60	57.74	r025	30/21	859903	0.37	15.63	12.74
r026	30/16	640655	3.27	26.96	23.06	r026	30/18	621194	0.60	18.73	15.60
r027	30/16	483131	2.49	32.47	29.10	r027	30/18	722126	0.56	11.22	8.00
r028	30/20	430367	12.06	47.09	43.43	r028	30/15	555117	1.64	26.91	24.95
r029	30/26	365148	11.46	44.39	40.65	r029	30/18	769308	0.55	5.99	3.93
Average Values:			8.51	41.16	37.61	Average Values:			0.90	18.82	15.83

Table 4.3: Experimental results obtained using the generalised DVS optimised scheduling approach for benchmark examples TG1 and TG2

resent improvements of 9.61% and 2.83%. Note that these reductions were obtained on benchmarks which do *not* show any power variations and so this optimisation feature of the proposed DVS algorithm stays unexploited. The achieved improvements are due to the fact that the proposed iterative GA based scheduling approach is able to explore a large space of potentially low-energy schedules, as opposed to the constructive list scheduling approach used in [64]. Regarding the computational times, Gruian *et al.* reported average times for the 30-node task graphs of 10s to 120s (on 440MHz, UltraSparc Iii, 256MB Workstation), while the proposed algorithm executes on average in 0.9s to 8.51s (on 750MHz, Pentium III PC, 128MB). These longer execution times of the constructive technique can be explained by the fact that the scheduling used in [64] has to re-schedule the tasks when no feasible schedule is found, i.e., the task priorities are re-adjusted. Further, the UltraSparc Workstation provides approximately 2–3 less computational power

compared to the PentiumIII PC according to SPECint performance evaluation [7]. However, the optimisation times indicate an advantage of the presented EE-GLSA technique.

In summary, this section has shown that significant improvements in terms of energy savings can be made by optimising the execution order of tasks and communications. In particular when compared to constructive techniques, the new EE-GLSA can achieve higher energy savings due to the effective exploration of the scheduling search space.

4.2 Optimisation of Task and Communication Mapping

Section 4.1 has shown that substantial energy savings (up to 17.32% were observed) can be achieved through an schedule optimisation that improves the execution order of tasks and communications not only towards performance goals, but additionally towards the effective utilisation of dynamic voltage scaling. Clearly, application mapping (Section 1.3.2) and activity scheduling (Section 1.3.3) are two heavily interrelated co-synthesis steps. For instance, mapping parallel tasks onto a single processing element would necessitate to execute these tasks one after the other (sequentially). On the other hand, mapping these tasks to different processing elements would allow to execute the tasks in parallel. This means that the mapping of tasks and communications has an important influence on the schedulability as well as on the utilisation of DVS and hence should be subject to optimisation from a timing and energy point of view. This section introduces a novel *two-step* approach that aims to improve the mapping towards these goals. Conceptually, the mapping approach *separates* the optimisation of task and communication mapping, i.e., both assignments are carried out in isolation of each other, as illustrated in Figure 4.1 (page 64). Correspondingly, this section is divided into two sections. Section 4.2.1 introduces a task mapping based on genetic algorithms that has been adapted to suit the particular problem of optimising the design for the effective exploitation of the DVS-PEs. Section 4.2.2 proposes a new method that extends the scheduling technique outlined in Section 4.1 to a combined optimisation of communication mapping and scheduling. Again, this technique aims at performance improvement as well as DVS utilisation. Using these techniques, the influence of mapping on the achievable energy savings through DVS is analysed in Section 4.2.3 for a set of benchmark experiments.

4.2.1 Genetic Task Mapping Algorithm

The task mapping step determines which PE carries out which task. Thereby, it determines the execution time and power dissipation of a task at nominal supply voltage and further the area requirement in terms of bytes or gates, whether a task is implemented as software or hardware. The goal of the mapping optimisation step is to distribute the tasks among the processing elements that form the distributed architecture, including the DVS-enabled PEs, such that the energy dissipation is minimised and feasible designs in terms of timing behaviour and area constraints are achieved. As mentioned in Section 2.4, task mapping has been intensively researched over the last decade. And similar to the scheduling problem, it belongs to the class of NP-hard problem [58]. That is, optimal solution for realistic problem sizes may only be found through extremely computational expensive processes. One effective way to address this problem is the usage of genetic algorithms for task mapping [43, 45]. Nevertheless, previous approaches did not consider the presence of DVS-PEs. In addition, as opposed to the mapping approach introduced in [43, 45], where solely the mapping of tasks is optimised by a GA, the approach presented in this thesis uses two independent GAs to find improved solutions for the mapping of communications and tasks. This section focuses on task mapping.

In GA based task mapping approaches, solution candidates (potential mappings) are encoded into mapping strings, as shown in Figure 4.10. Each gene in these strings describes a mapping of a task to a processing element. For instance, task τ_4 in Figure 4.10 is mapped to PE0. Similarly to the genetic algorithm used for the schedule optimisation (Section 4.1.1), the genetic task mapping algorithm (EE-GTMA) evolves a population of possible mapping solutions towards high quality implementations. The genetic algorithm that is employed to work on the task mapping strings is described in Figure 4.11. As typical in all genetic algorithms, the EE-GTMA applies ranking, selection, crossover, mutation, and offspring insertion in order to evolve an initial solution pool. The key feature of the EE-GTMA, however, is the invocation of the genetic list scheduling algorithm (EE-GLSA) for each mapping candidate, which allows to calculate parts of the fitness function that guides the optimisation. More precisely, the scheduling fitness F_S is used within the




Figure 4.10: Task mapping string describing the mapping of five tasks to an architecture

task mapping fitness F_M , as shown in the following equation:

$$F_M = F_S \cdot \prod_{\pi \in \mathcal{P}} AP_{\pi} \quad (4.4)$$

$$AP_{\pi} = \begin{cases} 1 & \text{if } AA_{\pi} \geq SA_{\pi} \\ k \cdot \left(\frac{SA_{\pi}}{AA_{\pi}} - 1 \right) + 1 & \text{otherwise} \end{cases} \quad (4.5)$$

where F_S is the schedule fitness (Equation (4.1), Section 4.1.3) based on the DVS reduced energy dissipation and a time penalty as outlined in Section 4.1.3. AP_{π} assigns an area penalty for each PE exceeding its area constraints as given in Equation (4.5). The used area is denoted SA_{π} and the maximal available area is represented by AA_{π} (either as memory or silicon area depending on the implementation in SW or HW). If the available area AA_{π} is not exceeded, it is not necessary to assign an area violation penalty for the particular processing element π , hence, F_S is multiplied by one. On the other hand, if the area constraint is exceeded, the used area SA_{π} and the available area AA_{π} are related and multiplied by a constant k , which allows to adjust the aggressiveness of the penalty. Through extensive experimentations, a value of 0.02 was found to be a good choice for the constant k , which was sufficiently high to avoid infeasible results at the end of the mapping optimisation. However, this value for k is still low enough to allow infeasible solutions to survive sometimes in order to increase the population diversity and avoid a premature convergence of the GA towards solutions of unnecessary low quality. In this

Algorithm: EE-GTMA

Input: - task graph TG
 - technology library (execution times, power dissipations)
 - allocated architecture

Output: - timing, area, and energy optimised mapping

- 01: **Initialisation:** Create initial population pool P_m of mapping strings, generated randomly
- 02: **Perform Mapping:** Generates, for each member of the solution pool, a mapping based on the corresponding mapping string. Specifies the task properties such as execution time, power dissipation, etc.
- 03: **Invoke EE-GLSA:** Invoke the schedule optimisation to, determine a suitable and energy efficient schedule for the current task mapping.
- 04: **Assign Fitness:** Compute fitness of each individual in the population pool.
 - a) Calculate area penalty
 - b) Derive fitness based on area penalty and the schedule fitness.
- 05: **Termination:** If no improved individual (improvement $> 1\%$) has been produced for 10 generations, then terminate. Otherwise, continue.
- 06: **Ranking:** Individuals are ranked according to their fitness.
- 07: **Selection:** According to the size of the generational overlap select individuals for mating. High ranked individuals have a high probability to be selected.
- 08: **Mating:** Produce two-point crossover between a pair of selected individuals.
- 09: **Mutation:** Randomly change genes of individuals using a dynamic mutation probability scheme, with exponential decreasing probability during run-time.
- 10: **Offspring insertion:** Exchange low ranked individuals by newly produced individuals with respect to the size of the generational overlap. Continue with step 02.

Figure 4.11: Proposed EE-GTMA approach for energy-efficient task mappings

way, it is possible to stimulate the placement of functionality onto the distributed PEs such that energy is minimised, while timing and area constraints are respected. The parameters of the GA for the task mapping were set as follows: The population size was set to 50, the minimal dynamic mutation probability was adjusted to 5%, the generational gap comprises 20%, and the initial population pool was filled with random mappings.

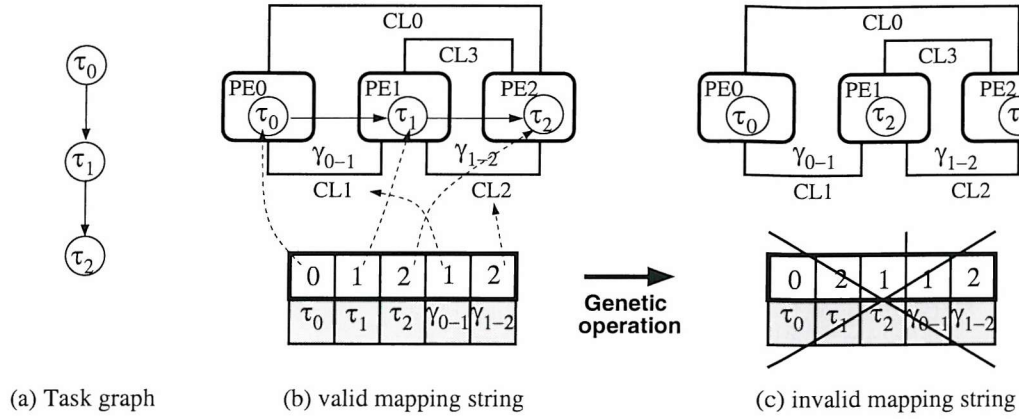


Figure 4.12: Combined optimisation of task and communication mapping

4.2.2 Combined Scheduling and Communication Mapping

The mapping approach in the previous section is used for the assignment of tasks to the processing elements only. However, communication issues have a great impact on the timing behaviour of the application and therefore should be considered carefully during the design space exploration [46, 79, 112], in order to find energy-efficient systems. One important decision that has been taken in this regard was the separation of communication mapping from the task mapping within the synthesis approach, i.e., communication and task mapping are carried out in two separate optimisation steps (see Figure 4.1). The following example illustrates the reasons behind this decision.

Example

The three tasks and two communications of the task graph shown in Figure 4.12(a) need to be mapped onto a target architecture consisting of three PEs, connected by four CLs. The string shown in Figure 4.12(b) combines tasks mapping (τ_0 , τ_1 , τ_2) and communication mapping (γ_{0-1} , γ_{1-2}) into one representation. This mapping represents a valid solution since communication γ_{0-1} between the tasks τ_0 and τ_1 is assigned to CL1, which connects the processing elements that accommodate tasks τ_0 and τ_1 . In a similar way, the communication γ_{1-2} is mapped onto a communication link (CL0) that interconnects the processing elements of task τ_1 and task τ_2 . Let us consider a certain genetic operation which transforms the valid mapping string shown in Figure 4.12(b) to the one in Figure 4.12(c). It can be observed that the mapping of the tasks τ_1 and τ_2 has been modified, while the communication mapping stays unchanged. A quick check upon this mapping indicates that

this assignment of activities represents an invalid solution. Consider, for example, the communication γ_{0-1} between task τ_0 and τ_1 . Although the tasks are mapped onto PE0 and PE2, which are solely connected through CL0, the communication is mapped to CL1. Hence, this mapping is invalid (if it is considered that only direct communications are allowed, i.e., communications without routing over intermediate PEs). Due to this reason, a combined task and communication mapping approach would produce a *high number of invalid solutions* during the GA based optimisation, which, in turn, would have a negative effect on the convergence of the population towards high quality solutions. To overcome this problem, the proposed mapping approach explicitly separates task and communication mapping. In précis, the introduced communication mapping technique is carried out in conjunction with the scheduling optimisation within the innermost loop of the proposed synthesis approach (see Figure 4.1). Hence, for each task mapping candidate, the scheduling and communication mapping are simultaneously optimised. In this way it is possible to avoid invalid solutions, since all possible mappings of communication activities onto the communication links are statically known for a particular task mapping. To clarify this consider again Figure 4.13. If the tasks τ_0 and τ_1 , for instance, are mapped to PE1 and PE2, respectively, then the communication γ_{0-1} can only be mapped onto CL2 or CL3. Hence, during the optimisation of the communication mapping it is possible to restrict the search to such feasible communication links. The proposed communication mapping, described next, takes advantage of this information to ensure that only valid solutions are produced. Thereby, the search space is restricted to structurally viable solutions only, decreasing significantly the synthesis run-time.

As mentioned above, the presented communication mapping optimisation is carried out in parallel with the schedule optimisation. To explain this strategy consider the extended string representation shown in Figure 4.13, which encodes both a possible schedule and a communication mapping candidate. It can be observed that this string is divided into priority and communication mapping genes. A list scheduler determines an execution order based on the encoded priorities, whilst the mapping of communication activities onto the interconnecting links is given by the communication mapping genes. The combined string representation allows the concurrent optimisation of priorities and communication mappings, using a single genetic algorithm. However, it necessitates a specialised genetic mutation, which operates on the two string parts without interference, i.e., random modifications on priorities need to be considered differently than the modifi-

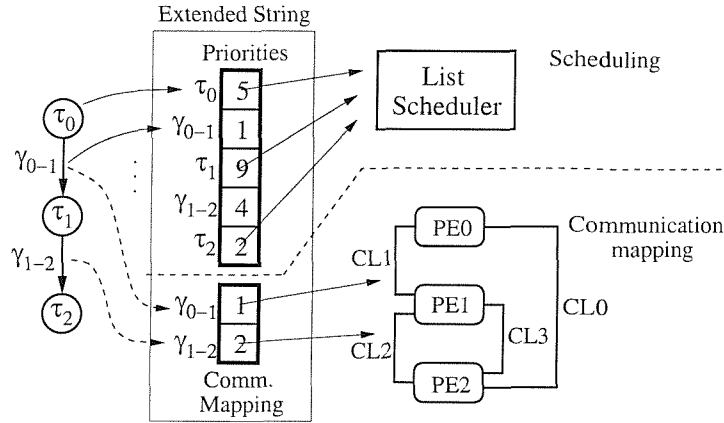


Figure 4.13: A combined priority and communication mapping string

cations on the communication mappings. On the other hand, standard crossover methods can be directly applied without worrying about feasibility, since the crossover between two strings maintains spatial locality, i.e., priority genes are not mixed with the mapping strings. Nevertheless, during initialisation and mutation of the communication mapping genes only valid values, which result in feasible communication mapping solutions, are allowed. This is not hard to achieve due to the fact that the task mapping precedes the communication mapping (scheduling and communication mapping are carried out in the innermost loop of the synthesis). Thereby, for every communicating pair of tasks the viable CLs are unambiguously specified. Therefore, it is possible to generate random initial chromosomes (random in the sense that a random choice is taken among the possible CLs) that assure proper communication mappings. Similarly, the mutation operator chooses randomly among the valid possibilities only. In order to keep the optimisation time low, the communication mapping string is dynamically adapted to the particular task mapping, as the number of inter-PE communications can change for each potential task mapping. Thereby, the amount of genes in the mapping string is kept at minimum, avoiding needlessly high synthesis times. Of course, the valid values of each gene change also dynamically in accordance to the task mapping. Note that the presented communication mapping optimisation improves both the timing behaviour as well as the power consumption, since the guiding fitness (Equation (4.1)) accounts for both.

The pseudo code of the combined scheduling and communication mapping algorithm, called EE-GLSCMA, is shown in Figure 4.14, an extended version of the EE-GLSA algorithm described in Figure 4.8 (page 76). The following modifications can be identified: In Step 01, the proposed combined optimisation of communication mapping and scheduling

Algorithm: EE-GLSCMA

Input: - task graph TG
 - task mapping and, correspondingly, the execution properties

Output: - timing and energy optimised schedule
 - timing and energy optimised communication mapping

- 01: **Find Feasible Communication Mappings:** Identify possible mappings of communication events to links, depending on the given task mapping.
- 02: **Initialisation:** Create initial population pool P of combined priority and communication mapping strings. Starting priorities are half randomly generated and half based on mobility, while initial, feasible communication mappings are randomly created.
- 03: **Perform List Scheduling:** Generates, for each member of the solution pool, a schedule based on the corresponding priority string.
 - a) Map communication to links
 - b) Assign task priorities from the property string
 - c) Invoke list scheduler without hole filling
- 04: **Perform Voltage Scaling:** Invoke the generalised DVS technique, calculating supply voltages for each task executed on a DVS-PE. This is done under the consideration of the individual power dissipation of tasks.
- 05: **Assign Fitness:** Compute fitness of each individual in the population pool.
 - a) Calculate timing penalty
 - b) Calculate energy based on the supply voltages
 - c) Derive fitness based on energy and timing penalty
- 06: **Termination:** If no improved individual (improvement $> 1\%$) has been produces for 10 generations, then terminate. Otherwise, continue with step 07.
- 07: **Ranking:** Individuals are ranked according to their fitness.
- 08: **Selection:** According to the size of the generational overlap select individuals for mating. High ranked individuals have a high probability to be selected.
- 09: **Mating:** Produce two-point crossover between a pair of selected individuals.
- 10: **Mutation:** Randomly change genes of individuals using a dynamic mutation probability scheme, with exponential decreasing probability during run-time. Mutation of communication mappings are randomly selected out of feasible assignments, depending on the task mapping.
- 11: **Offspring insertion:** Exchange low ranked individuals by newly produced individuals with respect to the size of the generational overlap. Continue with step 03.

Figure 4.14: Proposed EE-GLSCMA approach for combined optimisation of energy-efficient schedules and communication mappings

finds, for all data transfers between tasks, the feasible communication links. In addition to the initial priorities, Step 02 needs to generate random yet feasible communication mappings for the start of the optimisation. During Step 03, which performs the list schedul-

ing, the communication properties are calculated, in order to allow for the consideration of contention over the CLs. Finally, in Step 10, a specialised mutation is carried out, in order to avoid creation of invalid strings.

4.2.3 Experimental Results: Mapping Optimisation

The experimental results in this section analyse the effect of task and communication mapping on the achievable energy reductions through DVS. To evaluate this influence three basic concepts are compared:

CSE (Constructive Scheduling with fixed power DVS) Within this approach, DVS is tackled under the fixed power model (i.e., no power variations are considered). Scheduling is carried out using a mobility-based constructive list scheduling, as the one used in [89]. Communication mapping is based on a heuristic method that assigns communications to the first available CL. The task mapping is carried out using the genetic task mapping algorithm of Section 4.2.1.

DLSP (Dynamic Level Scheduling with power variation DVS) DLSP is based on a constructive list scheduling which dynamically re-calculates task priorities during the schedule construction [132]. The priorities of tasks are given by the dynamic level, which depends on the longest path of the activity and the earliest start time. The priorities are used for a combined task mapping and scheduling. Thus, this approach is thoroughly of constructive nature. This scheduling and mapping approach corresponds to the one used in [19]. The so produced schedules are scaled using the PV-DVS algorithm introduced in Chapter 3.

ICMSP (Iterative Combined Mapping and Scheduling with PV-DVS) This approach corresponds to the techniques and algorithms proposed in this thesis. That is, voltage scaling is performed by PV-DVS (Section 3.2.1), which considers the power variation model. Scheduling and communication mapping are based on the combined genetic algorithm EE-GLSCMA (Section 4.2.2). The task mapping is optimised using the genetic task mapping algorithm EE-GTMA (Section 4.2.1).

To ease the following discussion, Figure 4.15 provides a short overview of these optimisation concepts for reference purpose. All experimental results presented in this section are

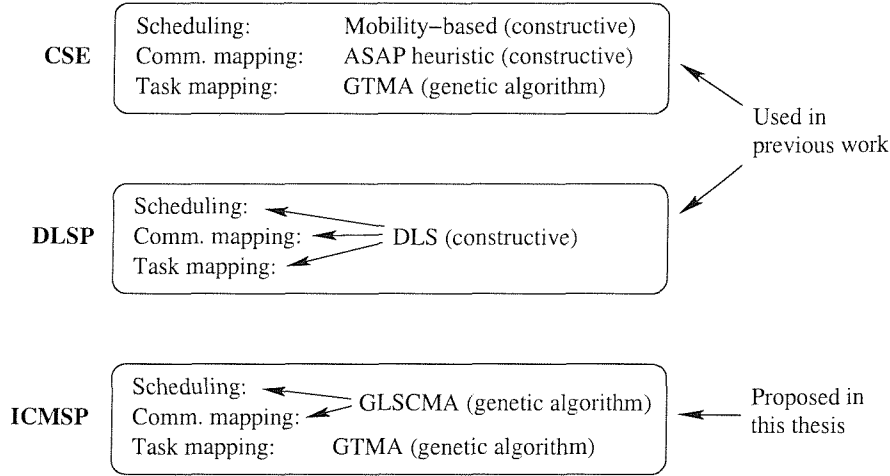


Figure 4.15: Three scheduling and mapping concepts

based on the same four benchmark sets (a–d) that have already been introduced in Sections 3.3 and 4.1. The presented results were obtained by running the optimisation process ten times and averaging the outcomes. The experiments are subdivided into two sections. Section 4.2.3.1 compares the CSE approach with ICMSP technique, while Section 4.2.3.2 assesses the optimisation potentials of DLSP with respect to ICMSP.

4.2.3.1 Comparison between CSE and ICMSP

The first experiments give a comparison between the CSE approach and the ICMSP approach. Table 4.4 shows this comparison for the benchmark sets *tgff* and *Hou*. The first column gives the benchmark names. The second column shows the nominal energy dissipations. These values were produced using the genetic task mapping algorithm (EE-GTMA) in combination with the genetic scheduling and communication mapping algorithm (EE-GLSCMA), however, *without* using dynamic voltage scaling. Thus, all tasks are executed at nominal supply voltage. The third column gives the computational times for these optimisations. The nominal energy values are used as base for a comparison of the CSE and ICMSP approaches. The achieved results of the CSE approach are shown in Columns 4–6, while the outcomes of ICMSP are given in Columns 7–9. The last column provides the achieved reduction factors, comparing CSE and ICMSP. Consider, for instance, the benchmark example *tgff4*, which has a nominal energy dissipation of $63924\mu J$. This solution was found after optimising the application mapping and scheduling for 24.15s. The energy dissipation of the same example is reduced to $15743\mu J$ using

Example	NO-DVS		CSE			ICMSP (proposed)			
	Energy Dissip. (μJ)	CPU time (s)	Energy Dissip. (μJ)	CPU time (s)	Reduc. (%)	Energy Dissip. (μJ)	CPU time (s)	Reduc. (%)	Reduc. Factor
tgff1*	333	3.11	116	1.91	65.23	92	12.14	72.41	1.11
tgff2	709747	24.10	625970	13.34	11.80	445532	47.86	37.23	3.15
tgff3	298991	69.46	225433	39.68	24.60	109351	2437.98	63.43	2.58
tgff4	63924	24.15	15743	12.15	75.37	10817	290.10	83.08	1.10
tgff4_t	49807	22.93	20275	11.83	59.29	18487	226.93	62.88	1.06
tgff4_fixed	59294	20.32	18860	11.66	68.19	10621	299.45	82.09	1.20
tgff5	568210	64.42	426614	41.23	24.92	233063	904.99	58.98	2.37
tgff6	24685	19.97	7298	11.66	70.44	3799	221.99	84.61	1.2
tgff7	1491203	10.29	1169258	5.55	21.59	1058346	41.75	29.03	1.34
tgff8	525250	15.52	182894	8.49	65.18	136057	46.92	74.1	1.35
tgff9*	600428	9.01	358087	4.63	40.36	323158	45.28	46.18	1.14
tgff10	9417	7.45	8531	3.98	9.41	7193	17.58	23.62	2.51
tgff11	2858919	26.87	2400940	14.48	16.02	2229397	97.6	22.02	1.37
tgff12	174440	56.36	90087	37.61	48.36	58404	1328.52	66.52	1.38
tgff13	927704	60.97	511019	32.56	44.92	328377	853.42	64.6	1.44
tgff14	7723	23.29	7578	14.39	1.88	6693	69.01	13.34	7.11
tgff15	20017	86.85	17948	54.39	10.34	16938	916.98	15.38	1.49
tgff16	2984716	34.66	2177495	24.64	27.05	2141352	197.41	28.26	1.04
tgff17	16237	41.97	11417	26.80	29.69	8220	308.54	49.38	1.66
tgff18	1518517	4.17	1248236	2.80	17.80	1066350	9.71	29.78	1.67
tgff19	3431	5.91	2176	4.12	36.59	1907	18.31	44.41	1.21
tgff20*	18621	12.41	7286	7.18	60.87	4646	92.24	75.05	1.23
tgff21	2182722	121.95	1543090	59.71	29.30	1352422	1665.04	38.04	1.3
tgff22	894765	301.48	691269	172.38	22.74	456021	2240.57	49.03	2.16
tgff23*	5519226	147.33	3261600	87.85	40.90	2129198	14050.26	61.42	1.5
tgff24	720861	151.80	302288	98.07	58.07	200328	2199.39	72.21	1.24
tgff25	3232360	74.20	2555077	44.36	20.95	2328983	1664.63	27.95	1.33
Hou*	11816	10.57	10704	11.43	9.41	6708	163.78	43.23	4.59
Hou_clust.*	12766	1.58	10145	1.97	20.53	7879	3.42	38.28	1.86

Table 4.4: Mapping optimisation with and without DVS optimised scheduling using tgff and hou benchmarks

the CSE approach, a reduction by 75.37%. This has been achieved in smaller run-time, although the CSE additionally account for DVS. The lower optimisation time can be explained by the fact that a quick constructive scheduling technique is used. Furthermore, EVEN-DVS (fixed power model) does not significantly increase the optimisation time since it can be performed with linear computational complexity. Nevertheless, it is possible to further increase the energy savings using the ICMSP approach, which considers the power variation model as well as an iterative optimisation of the scheduling and mapping towards DVS.

Using ICMSP the energy consumption of example tgff4 is reduced to 10817 μJ , a reduction by 83.08%. Due to the iterative scheduling and the quadratic complexity of PV-DVS, the optimisation time increases to 290.10s. Comparing the achieved reductions of the constructive-based CSE approach (75.37%) with the introduced ICMSP approach (83.08%), a reduction factor of $83.08/75.37 = 1.10$ can be calculated. Observing the remaining reduction factors given in Table 4.4, it can be seen that ICMSP (based on the

technique proposed in this thesis) always results in lower energy dissipations than the CSE approach (based on traditional scheduling techniques and the fixed power model [89]). These reductions are mainly achieved due to two reasons:

- (a) The schedule solution space can be more thoroughly search by an iterative scheduling technique.
- (b) Considering the power variations allows a more accurate energy estimation to guide the optimisation process.

The run-times for the CSE technique varied between 1.91s (tgff1) and 172.38s (tgff22) for task graphs with up to 100 nodes. The ICMSP approach optimised all the examples in 3.42s (hou_clust) to 14050.26s (tgff22). Clearly, a trade-off between run-time and quality.

Timing Behaviour Improvements

In addition to the schedule discussions given in Section 4.1, the following observations about the scheduling optimisation can be made. The scheduling optimisation does not only significantly reduce the dissipated energy, but also improves the timing behaviour, leading to feasible implementations where constructive techniques might fail. This is of great importance since high quality solutions are likely to be found in design space regions where infeasible and feasible solutions are spatially placed closely together. Making a wrong decision might involve a more costly implementation of the system. To clarify this, consider the results obtained with the benchmark set TG1 from Gruian *et al.* [64], as shown in Table 4.5. Scheduling the system tasks based on a constructive list scheduling heuristic (mobility-driven) produces a single solution, which might be feasible or infeasible. Consider, for example, benchmark r000. In the case of this benchmark the constructive scheduling attempt fails and the implementation is marked infeasible (Column 4, "unsolved"). Thus, making it necessary to increase the performance of the allocated system for the given mapping. On the other hand, the proposed iterative GA-based list scheduling technique is able to improve infeasible schedules by providing feedback to the optimisation process and therefore feasible schedules might be found, as in the case of the task graph example r000 (Column 7). This effect is likely to appear in the presence of tight deadline specifications, such as the benchmark set TG1. It can be observed that for 18 out of 30 examples no feasible mapping could be found when using a mobility-

Example	NO-DVS		CSE			ICMSP (proposed)			
	Energy Dissip.	CPU time (s)	Energy Dissip.	CPU time (s)	Reduction (%)	Energy Dissip.	CPU time (s)	Reduction (%)	Reduction Factor
r000	798700	53.87	unsolved	18.60	n/a	586806	194.86	26.53	n/a
r001	759500	56.16	592674	13.87	21.97	399839	804.73	47.35	2.16
r002	744800	55.64	unsolved	16.51	n/a	551944	189.97	25.89	n/a
r003	994700	27.76	711887	15.98	28.43	554171	769.58	44.29	1.56
r004	886900	54.00	unsolved	19.97	n/a	566263	360.58	36.15	n/a
r005	744800	54.94	465853	16.75	37.45	373677	1596.67	49.83	1.33
r006	901600	36.88	unsolved	17.55	n/a	589469	827.22	34.62	n/a
r007	837900	55.20	unsolved	20.20	n/a	565731	269.07	32.48	n/a
r008	862400	30.63	unsolved	19.25	n/a	635426	207.46	26.32	n/a
r009	681100	53.24	424723	14.99	37.64	311751	1535.28	54.23	1.44
r010	705600	55.96	464257	17.88	34.20	325572	1155.62	53.86	1.57
r011	803600	32.11	558165	17.98	30.54	432977	421.37	46.12	1.51
r012	994700	49.54	unsolved	17.86	n/a	703960	310.65	29.23	n/a
r013	818300	45.58	unsolved	21.36	n/a	546724	315.27	33.19	n/a
r014	872200	56.16	618498	17.99	29.09	467035	2324.70	46.45	1.60
r015	813400	54.97	501944	17.12	38.29	427297	2891.08	47.47	1.24
r016	744800	26.52	unsolved	20.82	n/a	599931	141.81	19.45	n/a
r017	886900	56.54	unsolved	17.75	n/a	625210	121.10	29.51	n/a
r018	700700	54.81	unsolved	14.05	n/a	420155	581.39	40.04	n/a
r019	872200	28.29	unsolved	16.13	n/a	612864	172.78	29.73	n/a
r020	744800	27.92	unsolved	18.73	n/a	442314	357.80	40.61	n/a
r021	984900	53.81	726542	18.88	26.23	660399	2542.27	32.95	1.26
r022	818300	34.35	unsolved	20.66	n/a	467490	636.25	42.87	n/a
r023	798700	55.24	487590	17.64	38.95	353875	1455.71	55.69	1.43
r024	828100	54.25	unsolved	20.78	n/a	521388	455.70	37.04	n/a
r025	764400	55.28	378677	16.10	50.46	284444	2877.87	62.79	1.24
r026	877100	55.98	unsolved	16.02	n/a	624911	255.15	28.75	n/a
r027	715400	55.26	unsolved	14.18	n/a	483865	225.51	32.36	n/a
r028	813400	51.26	585257	18.38	28.05	443074	932.38	45.53	1.62
r029	656600	27.86	unsolved	19.08	n/a	439722	218.24	33.03	n/a

Table 4.5: Mapping optimisation of the benchmark set TG1 using NO-DVS (Nominal), EVEN-DVS, and PV-DVS

driven scheduling algorithm. Nevertheless, using the genetic list scheduling approach it was possible to find feasible mappings for all task graphs of TG1, with energy reductions of up to 62.79% (benchmark r025). Compared to the feasible mappings generated using the CSE technique, energy reductions of up to 7.11 times could be achieved. This higher quality results require longer optimisation times.

4.2.3.2 Comparison between DLSP and ICMSP

To further confirm the ability of the proposed approach to optimise the mapping as well as the scheduling towards an effective utilisation of DVS, it is compared next with mappings and schedules produced by a dynamic level scheduling algorithm (DLS) [132]. This scheduling technique can be considered to be more sophisticated than "simple" mobility list scheduling approaches, since it dynamically re-calculates the task priorities after each schedule decision. Table 4.6 gives this comparison between DLSP and ICMSP. The first

<i>Example</i>	NO-DVS	DLSP		ICMSP (proposed)			
	<i>Energy Dissip.</i>	<i>Energy Dissip.</i>	<i>Reduction (%)</i>	<i>Energy Dissip.</i>	<i>CPU time (s)</i>	<i>Reduction (%)</i>	<i>Reduction Factor</i>
fft1	29600	18154	38.67	14019	591.2	52.63	1.36
fft3	48000	36772	23.39	21452	1144.7	55.31	2.37
karp10	59400	47737	19.63	24055	755.1	59.50	3.07
meas	28300	25732	9.07	25732	8.5	9.07	1
qmf4	16000	12740	20.38	11097	202.3	30.64	1.50

Table 4.6: Comparison between DLS algorithm and the proposed scheduling and mapping approach using Bambha's benchmarks [19]

and the second column show the benchmark names and the nominal energy consumption, respectively. The third and fourth column represent the results obtained by the DLSP³, while the fifth and the sixth column correspond to the proposed ICMSP approach. Accordingly, column seven gives the reduction factors when comparing DLSP with ICMSP. Examining the results, it can be seen that ICMSP was able to reduce the dissipated energy of 4 out of 5 examples by up to 39.87% (59.50% – 19.63%, karp10) equivalent to a reduction factor of 3.07. Only in the case of meas the dissipated energy remained the constant. This can be explained by the highly serialised graph structure of this task graph, which constrains the scheduling order and hence the potential to optimise the schedule. Furthermore, the serialised execution restricts the application parallelism and the used DVS-PEs are of identical types, thus, the task mapping does not influence the scheduling results. Nevertheless, for fft1, fft3, karp10, and qmf4, which show more parallelism, the introduced synthesis approach (ICMSP) outperforms the DLS-based scheduling in terms of energy savings by up to 3.07 times (karp10: 19.63% compared to 59.50%). It should be noted, both DLSP as well as ICMSP use dynamic voltage scaling under the power variation model. Hence, the reported saving are solely introduced by an improve application mapping and activity scheduling. Certainly, due the constructive nature of DLS, it surpasses the iterative improvement ICMSP approach in terms of optimisation time. For all examples the DLS run-time is below an optimisation time of 1s, while the presented technique shows run-times between 8.5 and 1144.7s. However, an advantage of the proposed approach is the possibility to initialise the mappings and schedules with a pre-passed mapping and scheduling based on DLS. This ensures that solutions of at least DLS quality are obtained in the first generation of the GA optimisation, that is, in an "identical" optimisation time. Such an initial solution pool could then be further opti-

³The DLSP results are identical to the results presented in Table 3.5 (Section 3.3.1).

minised iteratively. In this way, the system designer can easily exploit the freedom to trade off synthesis time and solution quality.

Summary

This section has analysed the effect of application mapping under different constellations of scheduling and voltage scaling techniques. The conducted experiments have shown that substantial energy saving can be achieved by the iterative techniques introduced in this thesis, when compared to constructive scheduling and mapping approaches [132, 150], which have been used in previous work on energy minimisation through DVS [19, 89]. Furthermore, the experiments indicate an advantage of the proposed techniques over constructive approaches also in terms of schedulability in the presence of tasks with tight deadlines. The results reinforce the importance of a thorough exploration of the mapping and scheduling solution space. Clearly, the cost for these better results is higher computational time.

4.3 Optimisation of Allocation

Sections 4.1 and 4.2 have introduced techniques and algorithms for the optimisation of activity scheduling and application mapping, respectively. As outlined in Section 1.3, the overall goal of the co-synthesis process is to support the designer in finding the "most" suitable target architecture, i.e., the optimisation of the architecture allocation. In the proposed system-level design approach, this step is user-driven and thereby based on the knowledge and experience of the designer. It is assumed that the designer has predefined an architecture and the voltage scaling, scheduling, and mapping techniques help him to evaluate the quality of the allocation in terms of energy dissipation, cost, and feasibility. If an architecture does not prove to be satisfactory, the designer makes the necessary changes and evaluates again. In this way, it is also possible to trade off the different design goals and hence achieve multiple design alternatives. Similarly to the scheduling and mapping steps, the allocation of components has an influence on the usability of DVS. For example, it might be beneficial to reduce the workload on the system PEs by introducing a new PE or by re-allocating faster PEs. And thereby it could be possible to increase the deadline slacks in the system schedules and hence exploit them using

DVS, resulting in higher dynamic energy reductions, while increasing the product cost and the static power consumption. Clearly, this optimisation is based on the astuteness of the designer. The following experiment demonstrates the importance of the architecture allocation using a real-world example.

4.3.1 Experimental Results: Component Allocation

In order to assess the energy reduction capability of the proposed synthesis approach in terms of real-world applicability, this set of experiments are concerned with a real-life optical flow detection (OFD) application. This application is a sub-system of an autonomous model helicopter for traffic monitoring purpose [12, 64], and it consists of 32 tasks. In its current implementation the OFD algorithm runs on two ADSP-21061L DSPs, with an average current of $760mA$ at $3.3V$, hence, an average power dissipation of approximately $2.52W$. Due to the stringent power budget on board of the helicopter, including application critical sub-systems, it is necessary to keep the overall power dissipation under a certain limit. To reduce the power consumption to a minimal amount, DVS seems predestined, since the OFD algorithm shows an unnecessary high performance (12.5 frames of 78×120 pixels per second). However, a repetition rate of 6.25 frames per second is sufficient (at certain operational heights) to ensure correct flow detection, allowing to relax the system constraints. For experimental purpose, a hypothetical extension of the DSPs towards DVS capability is considered. It is taken into account that such an extension has an influence on the static power dissipated by the digital circuits, and therefore the static power is increased by 10%.

In the first part of this experiment the application constraints are kept fixed, i.e., the OFD algorithm needs to perform with a repetition rate of $12.5Hz$ (equivalent to the current implementation). In order to increase the usage of the application parallelism, three different architectures are used, which are built out of three to five DVS-DSPs and connected via a shared bus. In this way the OFD algorithm can be performed faster, i.e., additional system slack is introduced, which is exploitable by DVS. Table 4.7 reports on the findings. The first row represents the current implementation of the OFD algorithm, i.e., running on an architecture consisting of two DSPs without DVS technology. This implementation shows a total average power dissipation of $2.52W$. Now, consider the architectures with three to five DVS-enabled DSPs. In accordance to the number of al-

<i>Architecture</i>	<i>Static Power (W)</i>	<i>Dyn. Power (W)</i>	<i>Total Power (W)</i>	<i>CPU Time (s)</i>	<i>Reduction (%)</i>
2 DSPs (NO-DVS)	0.383	2.137	2.520	n/a	n/a
3 DVS-DSPs	0.574	1.371	1.945	148.3	22.8
4 DVS-DSPs	0.736	1.163	1.899	303.6	24.6
5 DVS-DSPs	0.898	1.132	2.030	381.9	19.4

Table 4.7: Increasing architectural parallelism to allow voltage scaling of the OFD algorithm

<i>Architecture</i>	<i>Static Power (W)</i>	<i>Dyn. Power (W)</i>	<i>Total Power (W)</i>	<i>CPU Time (s)</i>	<i>Reduction (%)</i>
2 DSPs (NO-DVS)	0.383	1.069	1.452	n/a	n/a
2 DVS-DSPs	0.413	0.394	0.807	783.5	44.4
3 DVS-DSPs	0.574	0.277	0.851	1107.2	41.4
4 DVS-DSPs	0.736	0.253	0.989	1393.4	31.9
5 DVS-DSPs	0.898	0.241	1.139	1634.7	21.6

Table 4.8: Relaxing the performance constraints of the OFD algorithm

located PEs, the static power consumption increases. However, the increased number of PEs allows to exploit the application parallelism more effectively, which, in turn, allows a faster execution of the OFD algorithm. This results in slack time, usable by the DVS-PEs to lower the dynamic power dissipation. As it can be observed from Table 4.7, all implementations using DVS-DSPs show a reduction in total power consumption (sum of static and dynamic power consumption) of up to 24.6%. Note that this reduction does *not* necessitate any performance degradation, while the cost of the system increases.

As mentioned before, the current implementation of the OFD algorithm shows an unnecessary high performance and it is therefore possible to relax the system constraints. Hence, in the following experiment, the repetition rate is reduced from 12.5Hz to 6.25Hz, i.e., an execution at half speed. This performance is still high enough to allow a correct flow detection. The results of this investigation are shown in Table 4.8. Observing the results shows that for all given architectures the power dissipation could be reduced significantly, by up to 44.4% when compared to a non-DVS implementation (first row in Table 4.8). It is interesting to observe that the power consumption of the nominal task execution is reduced as well. This is due to the fact that the two DSPs are considered to consume no dynamic power when no computations are performed (through clock-gating). Therefore, for half of the operational time the DSPs dissipate no dynamic




Figure 4.16: Nine different implementation possibilities of the OFD algorithm

power. However, among all implementation alternatives the architecture composed out of two DVS-PEs (second row in Table 4.8) is the favourite, since it achieves the highest energy savings at a low cost. Its favourite position comes from the fact that with each additionally allocated PE the static power consumption increases, while the achievable dynamic energy reductions decrease (caused by limited parallelism within the application). Again, this shows how important an accurate design space exploration is when synthesising DVS-enabled embedded systems.

The power dissipations of all synthesised OFD systems are additionally shown in Figure 4.16. The depicted energy values correspond to the total energy dissipation, i.e., static as well as dynamic power dissipation are considered. Observing this chart clearly indicates the advantage of an OFD system using a two DVS-DSP implementation at a reduced performance rate of 6.25Hz.

4.4 Concluding Remarks

This chapter has introduced new methods for the activity scheduling and the application mapping of energy-efficient distributed heterogeneous embedded systems. A novel two-step iterative co-synthesis approaches that separates task mapping from a combined

scheduling and communication mapping optimisation has been developed and applied to a set of benchmarks examples. By separating the two optimisation, the search space is pruned to structurally feasible solutions only, hence, potentially reducing the required synthesis times. It has been shown that the introduced methods not only achieve substantially higher energy savings when compared to constructive techniques (up to 39.87% for benchmark `karp10`), but additionally improve the schedulability. The practicality of the proposed scheduling and mapping techniques has been validated using a real-world OFD application, for which the energy consumption could be reduced by 24.6% without degradation in performance and by 44.4% when a reduction in performance can be tolerated. The cost for these solutions of improved quality are longer synthesis times.

Chapter 5

Energy-Efficient Multi-Mode Embedded Systems

Chapter 3 has shown how the power variation model helps to reduce the energy consumption of DVS-enabled embedded systems during voltage scaling. Furthermore, it was shown in Chapter 4 how an appropriate application mapping and activity scheduling enables an effective utilisation of DVS, in order to save energy. These techniques and algorithms have addressed the optimisation of embedded systems that perform one single application, such as a standalone MP3 decoder or an optical flow detection (OFD). Nevertheless, one key characteristic of many current and emerging embedded systems is their need to work across a set of different interacting applications and operational modes. For instance, modern mobile phones often integrate not solely the functionality required for communication purpose, but additionally integrate applications like digital cameras, games, and complex multimedia functions (MP3 players and video decoders) into the same single device. In this thesis, such embedded systems are referred to as *multi-mode embedded systems*. This chapter introduces a novel co-synthesis methodology for the design of energy-efficient multi-mode embedded systems. Starting from a specification model that captures both mode interaction and functionality, the developed co-synthesis technique maps the application under consideration of *mode execution probabilities* to a heterogeneous architecture with the aim to reduce the energy consumption through an appropriate resource sharing between tasks. The main principle by which the co-synthesis methodology achieves energy-efficiency is an implementation trade-off between the different operational modes. In general, modes with high execution probability should be

implemented more energy efficient (e.g., by moving more tasks to hardware) than modes with a low execution probability. Nonetheless, the implementation of modes is heavily interrelated, due to the fact that different modes share the same resources (architecture). For example, mapping an energy-critical task of a highly active mode into energy-efficient hardware might prohibit to implement a timing-critical task into hardware due to the restricted hardware area (see motivational example in Section 5.2). Clearly, a well balanced implementation of the operational modes is vital for a good system design. In addition, the co-synthesis approach further reduces the energy dissipation by adapting the system performance to the particular needs of the active mode, using dynamic voltage scaling as well as component shutdown. Furthermore, the voltage scaling method of Chapter 3 is extended to account for hardware PEs that are capable of executing tasks in parallel, however, rely on a single scalable supply voltage source.

The rest of this chapter is organised as follows. Preliminaries, regarding the specification and the architectural models are given in Section 5.1. In Section 5.2, the problem addressed in this chapter is motivated through illustrative examples. Section 5.3 surveys relevant previous work regarding multi-mode embedded system. A formulation of the problems at hand is provided in Section 5.4. Section 5.5 proposes a novel synthesis approach, in order to tackle the identified problems. Extensive experimental results, including a smart phone example, are given in Section 5.6. Finally, concluding remarks are expressed in Section 5.7.

5.1 Preliminaries

This section introduces the functional specification model (Section 5.1.1) and the architectural model (Section 5.1.2), which are fundamental to the co-synthesis framework outlined in this chapter.

5.1.1 Functional Specification of Multi-Mode Systems

The abstract model used for the specification of multi-mode embedded systems consists of two parts. In précis, it is based on a combination of finite state machines and task graphs, capturing both the interaction between different operational modes as well as the functionality of each individual mode. Structurally, each node in the finite state machine

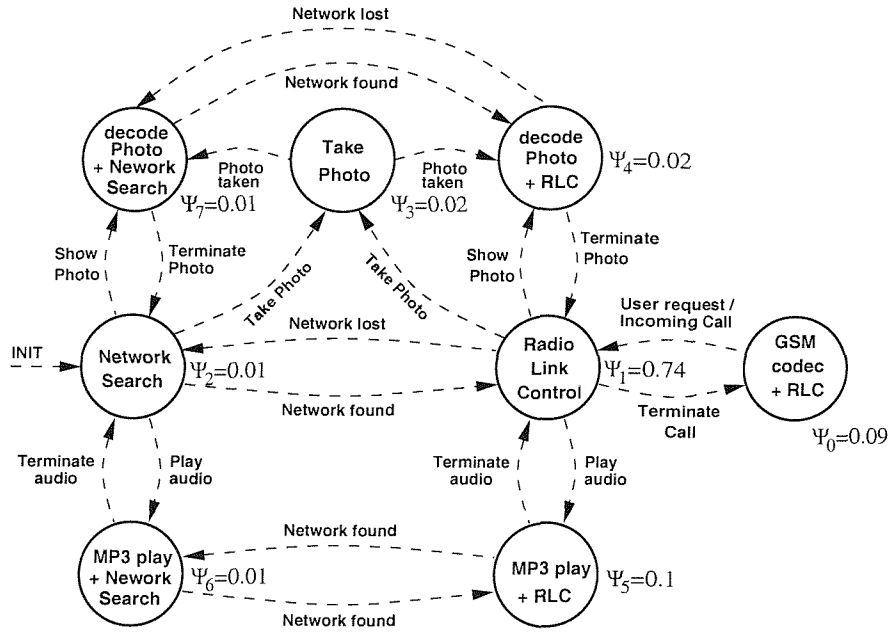


Figure 5.1: Example operational mode state machine of a smart phone

represents an operational mode and further contains the task graphs which are active during this mode. The following two sections introduce this model, which is henceforth referred to as operational mode state machine (OMSM).

Top-level Finite State Machine

In this work, it is considered that an application is given as a directed cyclic graph $\Upsilon(\Omega, \Theta)$, which represents a finite state machine. Within this top-level model, each node $O \in \Omega$ refers to an operational mode and each edge $T \in \Theta$ specifies a possible transition between two different modes. If the system undergoes a change from mode O_x to mode O_y , where $x \neq y$, the transition time t_T^{max} associated with the transition edge $T = (O_x, O_y)$ has to be met. At any given time there is only one active mode, i.e., the modes execute mutually exclusive. To exemplify the proposed model consider Figure 5.1. This figure shows the operational mode state machine for a smart phone example with eight different modes. A possible activation scenario could look like this: When switched on, the phone initialises into Network Search mode. The system stays in this mode until a suitable network has been found. Upon finding a network the phone undergoes a mode change to Radio Link Control (RLC). In this mode it maintains the connection to the network by handling cell hand-overs, radio link failure responses, and adaptive RF power control. An

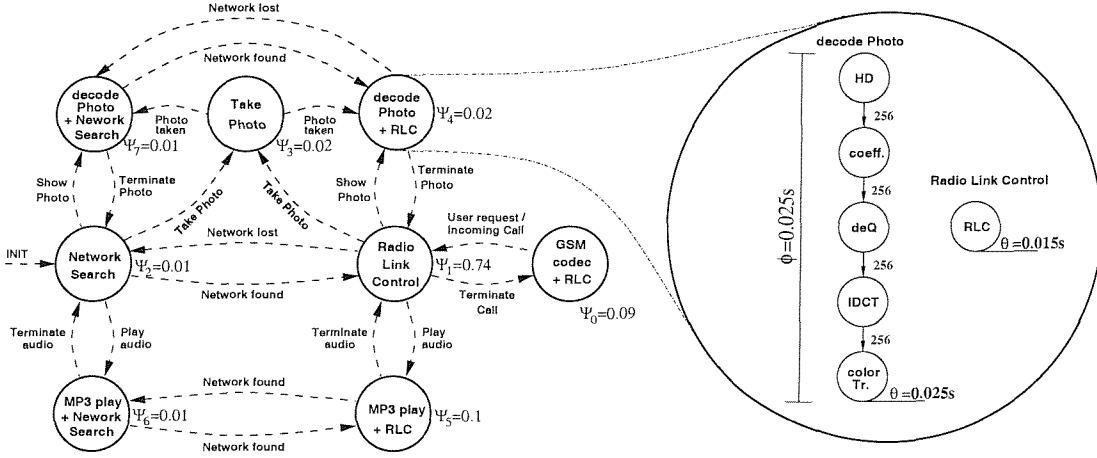


Figure 5.2: Relation between OMSM and individual task graph specifications

incoming phone call necessitates to switch the system into GSM codec + RLC mode. This mode is responsible for speech encoding and decoding, while simultaneously maintaining network connectivity. Similarly, the remaining modes have different functionalities and are activated upon mode change events. Such events originate upon user requests (e.g. MP3-player activation) or are initiated by the system itself (e.g. loss of network connection necessitates to switch the system into network search mode). Furthermore, based on the key observation that many multi-mode systems spend their operational time *unevenly* in each of the modes, an execution probability Ψ_O is associated with each operational mode O , i.e., it is known what percentage of the operational time the device spends in each mode. For instance, in accordance to the typical values given in Figure 5.1, the smart-phone stays 74% of this operational time in Radio Link Control (RLC) mode, 9% in GSM codec + RLC mode, and 1% in Network Search mode. The remaining 16% of the operation time are associated with the remaining modes. In practice the mode probabilities vary from user to user, depending on the personal usage behaviour. Nevertheless, it is possible to derive an average activation profile based on statistical information collected from several different users. Taking this information into account will prove to be important when designing systems with a prolonged battery lifetime.

Functional Specification of Individual Modes

The functional specification of each operational mode $O \in \Omega$ in the top-level finite state machine is expressed by a task graph $G_S^O(\mathcal{T}, \mathcal{C})$. This relation is shown in Figure 5.2. The task graph model was introduced in Chapter 1, Section 1.2.1. However, due to some par-

ticularities concerning multi-mode systems, the following outlines the exact model: Each node $\tau \in \mathcal{T}_O$ in a task graph represents a task, i.e., an atomic unit of functionality that needs to be executed without preemption. The level of granularity is coarse, i.e., tasks refer to functions such as Huffman decoder, de-quantizer, FFT, IDCT, etc. Therefore, each task is further associated with a task type $\eta \in \Gamma = \{HD, deQ, FFT, IDCT, \dots\}$. A distinctive feature of multi-mode systems is that task type sets $\Gamma^O \subseteq \Gamma$ of different modes $O \in \Omega$ can intersect, i.e., tasks of identical type can share the same hardware resource (inter-mode sharing). Resource sharing is also possible for multiple tasks of identical type that are found in a single mode (intra-mode sharing), however, due to task communalities among different modes, the chances to share resources are increased. Edges $\gamma \in \mathcal{C}$ in the task graph refer to precedence constraints and data dependencies between the computational tasks, i.e., if two tasks, τ_i and τ_j , are connected by an edge, then task τ_i must be finished and transfer data to task τ_j , before τ_j can be executed. A feasible implementation of a single mode O needs to respect all task deadlines θ , task graph period ϕ , and precedence relations.

5.1.2 Architectural Model and System Implementation

Similar to the techniques introduced in the Chapters 3 and 4, the proposed system-level synthesis approach targets distributed architectures that possibly consist of several heterogeneous processing elements (PEs), such as general-purpose processors (GPPs), ASIPs, ASICs, and FPGAs. These components are connected through an infrastructure of communication links (CLs). A directed graph $G_A(\mathcal{P}, \mathcal{L})$ captures such an architecture, where nodes $\pi \in \mathcal{P}$ and edges $\lambda \in \mathcal{L}$ denote PEs and CLs, respectively. Figure 5.3 shows an architecture example. Since each task might have multiple implementation alternatives, it can be potentially mapped onto several different PEs that are capable of performing this type of task. Tasks mapped to software-programmable components (i.e., GPP or ASIP) are placed into local memory. However, if a task is mapped to a hardware component (i.e., ASIC or FPGA), a core for this task type needs to be allocated. A feasible solution needs to obey the imposed area constraints, i.e., only a restricted number of cores can be implemented on hardware components. The subdivision of hardware components (ASICs and FPGAs) into hardware cores is shown in Figure 5.3. Each core is capable of performing a single task of type $\eta \in \Gamma$ at a time. Tasks assigned to GPPs or ASIPs (software tasks) need to be sequenced, whilst the tasks mapped onto FPGAs and ASICs (hardware tasks) can be

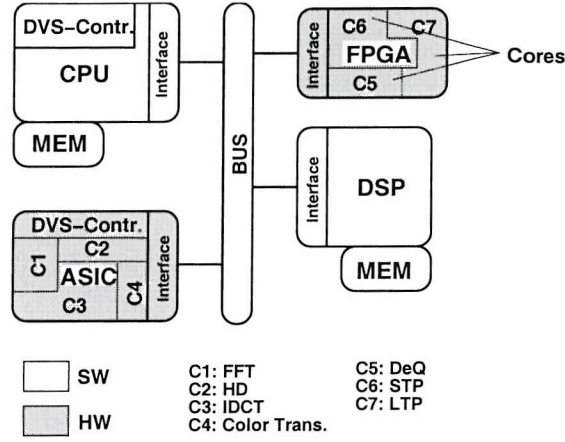


Figure 5.3: Distributed Architectural Model

performed in parallel if the necessary resources (cores) are not already engaged. However, contention between two or more tasks assigned to the same hardware core requires a sequential execution order, similar to software tasks. Cores implemented on FPGAs can be dynamically reconfigured during a mode change, involving a time overhead, which needs to respect the imposed maximal mode transition times t_T^{max} . Further, similar to the problems investigated in Chapter 3 and Chapter 4, PEs might feature dynamic voltage scaling to enable a trade-off between power consumption and performance that can be exploited during run-time. A set \mathcal{V}_π specifies the available discrete voltages of DVS-PE π . For such PEs a voltage schedule needs to be derived, in addition to a timing schedule. To implement a multi-mode application captured as OMSM, the tasks and communications of all operational modes need to be mapped onto the architecture, and a valid schedule for these activities $\varepsilon \in \mathcal{A}$, where $\mathcal{A} = \mathcal{T} \cup \mathcal{C}$, needs to be constructed. Further, for tasks mapped to DVS-enabled components an energy reducing voltage schedule has to be determined. According to these aspects, an implementation candidate can be expressed through four functions, which need to be derived for each operational mode $O \in \Omega$:

Task mapping: $M_\tau^O : \mathcal{T} \rightarrow \mathcal{P}$

Communication mapping: $M_\gamma^O : \mathcal{C} \rightarrow \mathcal{L}$

Timing schedule: $S_\varepsilon^O : \mathcal{A} \rightarrow \mathbb{R}_0^+$

Voltage schedule: $V_\tau^O : \mathcal{T}_{DVS} \rightarrow \mathcal{V}_\pi$

where M_τ^O and M_γ^O denote task and communication mapping, respectively, assigning tasks to PEs and communications to CLs. Activity start times are specified by the scheduling




Figure 5.4: Mode execution probabilities

function S_e^O , while V_τ^O defines the voltage schedule for all tasks $\tau \in \mathcal{T}_{DVS}$ mapped to DVS-PEs, where \mathcal{V}_π is the set of the possible discrete supply voltages of PE π . Clearly, the mappings as well as the corresponding schedules are defined for every mode separately, i.e., during the change from mode O_x to mode O_y , the execution of activities found in mode O_x are finished, and the activities of mode O_y are activated.

5.2 Motivational Examples

The aim of this section is to motivate the key ideas behind the new multi-mode co-synthesis, that is, the consideration of mode execution probabilities and multiple task type implementations. First, the influence of mapping in the context of multi-mode embedded systems with different mode execution probabilities is demonstrated. Second, it is illustrated that multiple task implementations can help to reduce the energy dissipation of multi-mode embedded systems.

Example: Mode Execution Probabilities

For simplicity, timing and communication issues are neglected in the following example. Consider the application shown in Figure 5.4(a), which consists of two operational modes, O_1 and O_2 , each specified by a task graph with three tasks. The system spends 10% of its operational time in mode O_1 and the remaining 90% in mode O_2 , i.e., the execution probabilities are given by $\Psi_1 = 0.1$ and $\Psi_2 = 0.9$. The specification needs to be mapped onto

a target architecture built of one general-purpose processor (PE1) and one ASIC (PE2), linked by a bus (CL1). Depending on the task mapping to either of the components, the execution properties of each task are shown in Table 5.1. It can be observed that all tasks

task type	PE1 (SW)		PE2 (HW)		
	exec. time (ms)	dyn. energy (mJ)	exec. time (ms)	dyn. energy (mJ)	area (mm ²)
A	20	10	2	0.010	24.0
B	28	14	2.2	0.012	30.0
C	32	16	1.6	0.023	27.5
D	26	13	3.1	0.047	24.5
E	30	15	1.8	0.015	21.0
F	24	14	2.2	0.032	28.0

Table 5.1: Task execution and implementation properties

are of different type, therefore, if a task is mapped to HW, a suitable core needs to be allocated explicitly for that task. Hence, in this particular example, no hardware sharing is considered. Each allocated core uses area on the hardware component that offers 60mm^2 , i.e., at most 2 cores can be allocated at the same time without violating the area constraint (see Table 5.1, Column 6). Note that although the two modes execute mutually exclusive, the task types implemented in hardware (HW cores) cannot be changed during run-time, since their implementation is static (non-reconfigurable ASIC); as opposed to software-programmable components. Consider the mapping shown in Figure 5.4(b), in which the highest energy consuming tasks (τ_3 and τ_5 , when implemented in software) are executed using a more energy-efficient hardware implementation. According to the task energy dissipations given in Table 5.1, the energy dissipation during modes O_1 and O_2 are $E_1 = 10\text{mJ} + 14\text{mJ} + 0.023\text{mJ} = 24.023\text{mJ}$ and $E_2 = 13\text{mJ} + 0.015\text{mJ} + 14\text{mJ} = 27.015\text{mJ}$. Neglecting the mode execution probabilities by assuming that both modes are active for even amounts of time (50% mode O_1 and 50% mode O_2) energy consumption can be calculated as $E_e = 0.5 \cdot 24.023\text{mJ} + 0.5 \cdot 27.015\text{mJ} = 25.519\text{mJ}$. Nevertheless, taking the real behaviour into account, mode O_1 is active for 10% of the operational time, i.e., its energy dissipation can then be calculated as $E_{r1} = 0.1 \cdot 24.023\text{mJ} = 2.4023\text{mJ}$. Similarly, mode O_2 is active 90% of the operational time, hence, its energy is given by $E_{r2} = 0.9 \cdot 27.015\text{mJ} = 24.3135\text{mJ}$. Based on both modes, the real energy dissipation results in $E_r = E_{r1} + E_{r2} = 26.7158\text{mJ}$. Now consider an alternative mapping for the same task graph, shown Figure 5.4(c). In this configuration tasks τ_5 and τ_6 , i.e., the most energy dissipating tasks of the highly active mode O_2 , use energy-efficient hard-

were implementations on PE2, while task τ_3 of the less active model O_1 is shifted into the software-programmable processor (PE1). According to this solution, the energy consumptions of modes O_1 and O_2 are given by $E_1 = 10mJ + 14mJ + 16mJ = 40mJ$ and $E_2 = 13mJ + 0.015mJ + 0.032mJ = 13.047mJ$. Considering the even execution of each mode (neglecting the execution probabilities), the even energy consumption can be calculated as $0.5 \cdot 40mJ + 0.5 \cdot 13.047mJ = 26.524mJ$. Note that this value is higher than the even energy of the first mapping ($E_e = 25.519mJ$). Thus, a co-synthesis approach that neglects the mode execution probabilities would optimise the system towards the first mapping. However, in real-life the modes are active for different amount of time and hence the real energy dissipation is given by $E_r = 0.1 \cdot 40mJ + 0.9 \cdot 13.047mJ = 15.7423mJ$. This is 41% lower compared to the first mapping ($E_r = 26.7158mJ$) shown in Figure 5.4(b), which is not optimised for an uneven task execution probability. Furthermore, the second task mapping allows to switch off PE2 and CL1 during mode O_1 , since all tasks of this mode are assigned to PE1. This results in a significant reduction of the static power, additionally increasing the energy savings.

Example: Multiple Task Type Implementations

An important characteristic of multi-mode systems is that tasks of the same type might be found in different modes, i.e., resources can be shared among the different modes in a time-multiplexed fashion. To increase the possibility of component shutdown, it might be necessary to implement the same task type multiple times, however, on different components. The following example, shown in Figure 5.5, clarifies this aspect. Here tasks τ_1 and τ_4 are of type A (see Figure 5.5(a)), allowing resource sharing between these tasks. The sharing is possible without contention due to the mutual exclusive execution of these tasks, since only one mode is active at a given time. In the first mapping, given in Figure 5.5(b), both tasks utilise the same HW core. However, implementing task τ_4 in software (additional task type A on PE1), as shown in Figure 5.5(c), allows to shut down PE2 and CL1 during the execution of mode O_2 . Hence, multiple implementations of task types can help to reduce power dissipation.

These two examples have demonstrated that it is essential to guide the synthesis process by: (a) an energy model that takes into account the mode execution probability as well as (b) allowing multiple task implementations.



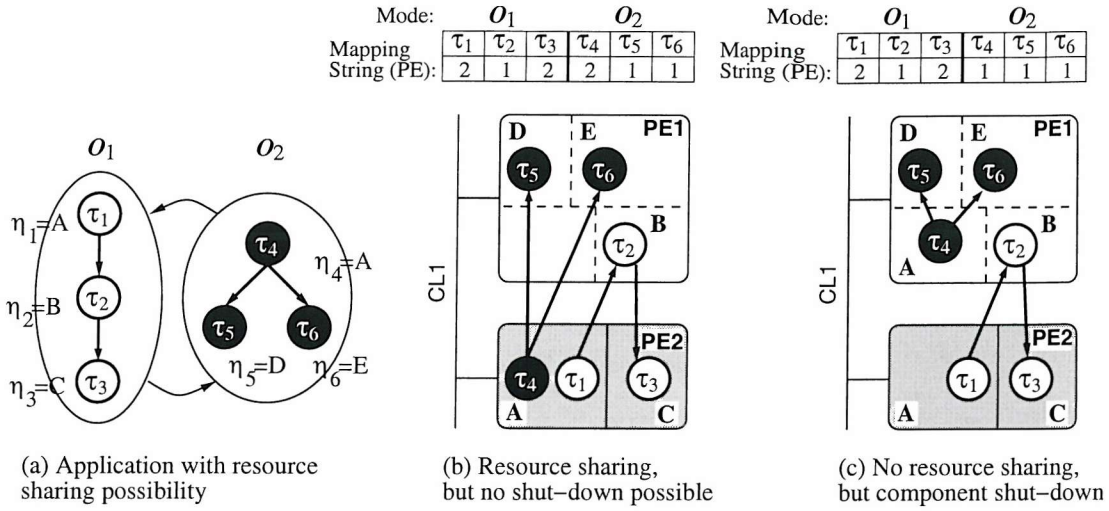


Figure 5.5: Multiple task type implementations

5.3 Previous Work

Over the last few year, numerous methodologies for the design of low power consuming embedded systems have been proposed, including approaches that leverage power management techniques, such as dynamic power management (DPM) and dynamic voltage scaling (DVS). Nevertheless, a crucial feature of many modern embedded systems is their capability to execute several different applications (multi-modes), which are integrated into a single device.

Approaches for the schedulability analysis of systems with several modes of operations can be found in the real-time research community [104, 129]. However, these approaches solely concentrate on scheduling aspects (i.e., they investigate if the mode change events fulfil the imposed timing constraints) and do not address implementation aspects. Three recent approaches have addressed various problems involved in the design of multi-mode embedded systems [76, 101, 130]. Shin *et al.* [130] proposed a schedulability-driven performance analysis technique for real-time multi-mode systems. They show that it is possible, through a sophisticated performance estimation, to identify timing critical tasks, which are active in different operational modes. This identification allows to improve the execution times of the most crucial tasks, in order to achieve system schedulability. In their work, the optimisation of the identified tasks is up to the designer. For example, reductions in the execution times can be made by handcrafted code tuning and outsourcing of core routines into hardware. Kalavade and Subrahmanyam [76] have

introduced a hardware/software partitioning approach for systems that perform multiple functions. Their technique classifies tasks, found within similar applications, into groups of task types. The implementation of frequently appearing task types is biased towards hardware. This can be intuitively justified by the fact that costly hardware implementations are shared across a set of applications, hence, exploiting the allocated hardware more effectively. Oh and Ha [101] address the problem in a slightly different way. Their co-synthesis framework for multi-mode systems is based on a combined scheduling and mapping technique for heterogeneous multiprocessor systems (HMP [99]). Taking a processor utilisation criterion into account, an allocation-controller selects the required processing elements such that the schedulability constraint is satisfied and the system cost is minimised. The main principle behind all three approaches is to consider the possibility of resource sharing, i.e., computational tasks of the same type, which can be found in different modes, utilise the same implementations. Thereby, multiple hardware implementations of the same task type are avoided, which, in turn, reduces the hardware cost. As opposed to these approaches, the work presented in this chapter addresses the design of low energy consuming multi-mode systems; hence, it differs in several aspects from the previous works. To the author's knowledge, there has been no prior work investigating the co-design problem of power minimisation in the context of multi-mode embedded systems. This chapter makes the following contributions:

- (a) The consideration of *mode execution probabilities* and their effect on the energy-efficiency of multi-mode embedded systems is analysed and demonstrated.
- (b) A co-design methodology for the design of energy-efficient multi-mode systems is presented. The proposed co-synthesis maps and schedules a system specification that captures both mode interaction and mode functionality onto a distributed heterogeneous architecture. Four mutation strategies are introduced that aid the GA based optimisation process in finding solutions of high quality by pushing the search into promising design space regions.
- (c) Dynamic voltage scaling is investigated in the context of multi-mode embedded systems. A transformation-based approach is used to tackle the problem of DVS on processing elements that execute different tasks in parallel, but offer only a single scalable supply voltage source.

5.4 Problem Formulation

The goal of the introduced co-synthesis is an energy-efficient and feasible implementation of application Υ , modelled as OMSM. This involves the derivation of the mapping and schedule functions, M_τ^O , M_γ^O , S_ε^O , and V_τ^O (outlined Section 5.1.2), under the consideration of static and dynamic power as well as mode execution probabilities. Although the technique and algorithms introduced in the Chapters 3 and 4 concentrate on the minimisation of the dynamic power, in multi-mode systems static power consumption can have a significant impact on the overall energy efficiency. The reasons for this are the different performance requirements of the various operational modes. For instance, the minimal performance requirements of the hardware architecture are imposed by the most computational intensive mode, i.e., the minimal allocated architecture has to provide enough computation power to execute this performance critical mode. However, the allocated architecture might be far too powerful for the execution of modes with low performance needs. Furthermore, low performance modes, such as the standby-mode of mobile phones (i.e., Radio Link Control), often account for the greatest portion of the system time. During such circumstances, the static energy dissipation of unnecessarily switched-on PEs and CLs can outweigh the dynamic energy consumption caused by tasks of a "lightweight" mode. Thus, switching-off the unneeded components becomes an important aspect particularly in multi-mode embedded systems. In accordance, an accurate estimation of the average power consumption of an implementation alternative should consider both static and dynamic power, and furthermore the mode execution probabilities. The average power consumption \bar{p} can be expressed using the following equation:

$$\bar{p} = \sum_{O \in \Omega} (P_O^{stat} + P_O^{dyn}) \cdot \Psi_O \quad (5.1)$$

where P_O^{stat} , P_O^{dyn} , and Ψ_O refer to the static power dissipation, the dynamic power dissipation, and the execution probability of mode O , respectively. The static and dynamic power consumptions are given as:

$$P_O^{stat} = \sum_{\xi \in \mathcal{K}_O} P^{stat}(\xi) \quad (5.2)$$

and

$$P_O^{dyn} = \left(\sum_{\varepsilon \in \mathcal{A}_O} E^{dyn}(\varepsilon) \right) \cdot \frac{1}{hp_O} \quad (5.3)$$

where $P^{stat}(\xi)$ refers to the static power consumption of a component ξ , which is found in the set of all active components $\mathcal{K}_O \subseteq (\mathcal{P} \cup \mathcal{L})$ of mode O . Further, \mathcal{A}_O and hp_O denote all activities and the hyper-period of mode O , respectively. With respect to the type of activities, the dynamic energy consumption $E^{dyn}(\varepsilon)$ can be calculated in the same way as introduced in Equation (4.2) of Section 4:

$$E^{dyn}(\varepsilon) = \begin{cases} P_{max}(\varepsilon) \cdot t_{min}(\varepsilon) \cdot \frac{V_{dd}^2(\varepsilon)}{V_{max}^2(\varepsilon)} & \text{if } \varepsilon \in \mathcal{T}_{DVS} \\ P_{max}(\varepsilon) \cdot t_{min}(\varepsilon) & \text{if } \varepsilon \in \mathcal{T} \setminus \mathcal{T}_{DVS} \\ P_C(\varepsilon) \cdot t_C(\varepsilon) & \text{if } \varepsilon \in \mathcal{C} \end{cases}$$

where P_{max} is the dynamic power consumption and t_{min} the execution time of tasks when executed at nominal supply voltage V_{max} . Tasks $\tau \in \mathcal{T}_{DVS}$ mapped to DVS-PEs can execute at a scaled supply voltage V_{dd} , resulting in a reduced energy consumption. Further, communications consume power P_C over a time t_C . The mode execution probabilities used in Equation (5.1) are either based on approximations or statistical information collected from several real users. In the case that statistical information is available from a set of different users U , the average execution probabilities $\bar{\Psi}_O$ of a single operational mode $O \in \Omega$ can be calculated.

The co-synthesis goal is to find a task mapping M_τ^O , a communication mapping M_γ^O , a starting time schedule S_ε^O , as well as a voltage schedule V_τ^O for each operational mode O , such that the total average power \bar{p} , given in Equation (5.1), is minimised. Furthermore, a feasible implementation candidate needs to fulfil the following requirements:

- (a) The mapping of tasks M_τ^O does not violate area constraints in terms of memory and hardware area, i.e., $(\sum_{\eta \in \Gamma_\pi} a_\eta) \leq a_\pi^{max}$, $\forall \pi \in \mathcal{P}$; where Γ_π is the set of all task types implemented on PE π , and a_η and a_π^{max} refer to the area used by task type η and the available area on PE π , respectively.
- (b) The timing schedule S_ε^O and the voltage schedule V_τ^O , based on task and communication mapping, do not exceed any task deadlines θ_τ or task graph repetition periods ϕ , therefore, $t_S(\tau) + t_{exe}(\tau) \leq \min(\theta_\tau, \phi)$, $\forall \tau \in \mathcal{T}$; where $t_S(\tau)$ and t_{exe} refer to task start time and task execution time (potentially based on voltage scaling).

- (c) The system reconfiguration time t_r between mode changes does not exceed the imposed maximal mode transition times t_T^{max} . Hence, $t_T \leq t_T^{max}$, $\forall T \in \Theta$ needs to be respected for all mode transitions.

5.5 Co-Synthesis of Energy-Efficient Multi-Mode Systems

Energy minimisation techniques for application mapping and activities scheduling of single mode embedded systems have already been introduced in Chapter 4. This section proposes new techniques for the co-synthesis of energy-efficient multi-mode embedded systems. Similar to the algorithms given in Chapter 4, the co-synthesis for multi-mode systems is based on two nested optimisation loops. The outer loop optimises task mapping and core allocation, while the inner loop is responsible for the combined optimisation of communication mapping and scheduling. While the communication mapping and scheduling in the multi-mode co-synthesis approach are based on the algorithms given in Section 4.2.2, this section reconsiders task mapping, hardware core allocation, and dynamic voltage scaling to suit the particular problems of multi-mode embedded systems. Hence, the remainder of this section is organised as follows. Section 5.5.1 outlines a new co-synthesis algorithm for multi-mode systems, concentrating on task mapping and four improvement strategies that aid to tackle the problem of multi-mode task mapping. Section 5.5.2 outlines a heuristic technique for hardware core allocation. Finally, Section 5.5.3 describes a transformation-based approach for dynamic voltage scaling of parallel execution task on single hardware components.

5.5.1 Multi-Mode Co-Synthesis Algorithm

The task mapping approach, which determines M_t for all modes of application Υ is an enhancement of the genetic task mapping algorithm (EE-GTMA) introduced in Section 4.2.1. These enhancements include the consideration of resource sharing, component shutdown, and mode transition issues. As outlined in Section 4.1.2, GAs optimise a population of individuals over several generations by imitating and applying the principles of natural selection. That is, the GA iteratively evolves new populations by mating (crossover) the fittest individuals (highest quality) of the current population pool until a certain convergence criterion is met. In addition to mating, mutation, i.e., the random change of genes

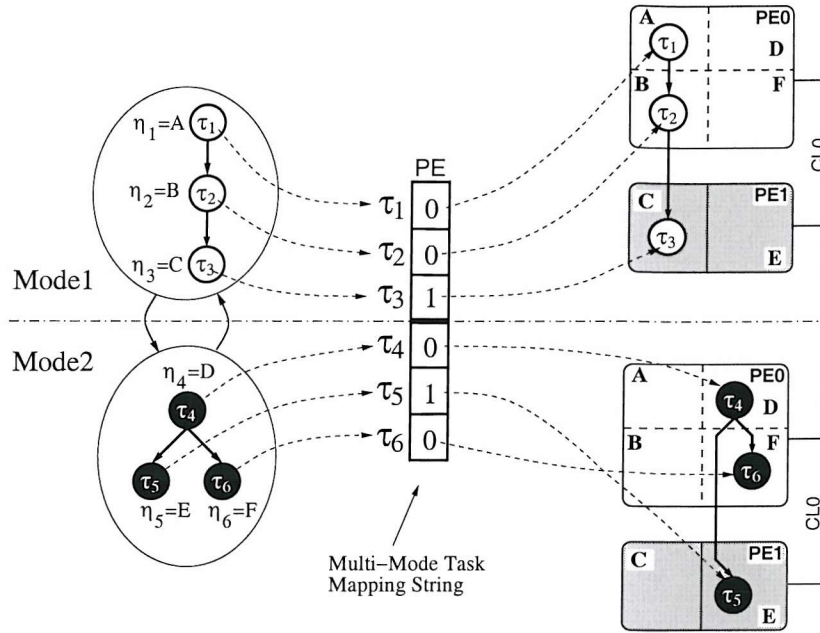


Figure 5.6: Task mapping string for multi-mode systems

in the genome (string), provides the opportunity to push the optimisation into unexplored search space regions. As opposed to the single mode task mapping strings of Section 4.2, task mapping strings for multi-mode specifications combine the mapping strings of each operational mode into one large task mapping string as shown in Figure 5.6. Within this string each number represents the PE to which the corresponding task is assigned.

The goal of the co-synthesis is to find a mapping of tasks that minimises the total power consumption and obeys the performance constraints. Figure 5.7 outlines the pseudo-code of the co-synthesis algorithm. Starting from an initial random population of multi-mode task mapping strings (line 1), the optimisation runs until the convergency criterion is met (line 2). The used criterion is based on the diversity in the current population and the number of elapsed iterations without producing any improved individual. To judge the quality of mapping candidates, i.e., the fitness which guides the genetic algorithm, it is necessary to estimate important design objectives, including static and dynamic power dissipation, area usage, and timing behaviour (lines 03–13). The following explains each of the required estimations. The hardware area depends on the allocated cores on each hardware component (ASIC or FPGA). Of course, for each task type mapped to hardware at least one core of this type needs to be allocated. However, if too many cores are placed onto a single ASIC or FPGA, the available area is exceeded and an area penalty is introduced (line 6). On the other hand, if multiple tasks of the same type are mapped to the

Algorithm: MULTI-MODE-SYN	
Input:	- OMSM (finite state machine + task graphs), Technology Library, Allocated Architecture
Output:	- Outer loop: Core Allocation, Task Mapping - Inner loop: Communication Mapping, Scheduling, Scaled Voltages
<pre> 01: Pop=CreateInitialPopulation // randomly 02: while(NoConvergence(Pop)) 03: forall map ∈ Pop 04: mob=ComputeMobilities(map) // ASAP & ALAP 05: cores=ImplementHWcores(map,mob) // (Section 5.5.2) 06: ap=CalcAreaPenalty(map,cores) 07: PstatPE=CalcStaticPowerPE 08: trp=CalcTransitionPenalty(cores) 09: forall mode ∈ Ω 10: CommMapping_Scheduling(mode) // inner loop (Section 4.2.2) 11: tp(mode)=CalcTimingPenalty(mode) 12: Pdyn(mode)=CalcDynPower(mode) // incl. DVS 13: PstatCL=CalcStaticPowerCL 14: FM=MappingFitness(Pdyn,tp,PstatCL,PstatPE,ap,trp) 15: ran=RankingIndividuals(FM) 16: mat=SelectedMatingIndividuals(ran) 17: TwoPointCrossover(mat) 18: OffspringInseration(Pop) 19: ShutdownImprovementMutation(Pop) 20: AreaImprovementMutation(Pop) 21: TimingImprovementMutation(Pop) 22: TransistionImprovementMutation(Pop) </pre>	

Figure 5.7: Pseudo Code: Multi-Mode Co-Synthesis

same hardware component and the hardware area is not violated, it is possible to implement cores multiple times (if helpful for the energy reduction). In the proposed approach, additional cores (line 5) are allocated for parallel tasks with low mobility (line 4), therefore, the chance to exploit application parallelism is increased. Clearly, from an energy point of view, this is also preferable, especially in the presence of DVS, where a decreased execution time can be exploited. Section 5.5.2 describes the core allocation in more detail. At this point it is possible to compute the static power consumption of the implementation (line 7), taking into account component shutdown between different modes. Components can be shut down during the execution of a certain mode whenever no tasks belonging to that mode are mapped onto these components, i.e., the component is vacant (for instance,

PE0 during execution of mode O_2 in Figure 5.5, page 112). Another important aspect is the reconfigurability of FPGAs which allows to exchange the implemented cores to suit the active mode. However, this reconfiguration during a mode change takes time, hence, a transition penalty is introduced if the maximal transition times are exceeded (line 8). Having determined the cores to be implemented (line 5), it is now also possible to schedule each mode of the application and to derive a feasible communication mapping (line 10). Since the modes are mutually exclusive, it is possible to employ a communication mapping and scheduling optimisation for a single mode system. The technique outlined in Section 4.2.2 is utilised for this purpose. If timing constraints are violated by the found schedule, a timing penalty is introduced (line 11). Furthermore, based on the communication mapping and scheduling, the dynamic power consumption of the application can be computed, taking into account DVS (line 12) if voltage-scalable components are present. Similarly to the shutdown of PEs, it is also possible to switch off a CL when no communications are mapped to that link (line 13), therefore, further reducing the static power consumption of the system. Based upon all estimated power consumptions and penalties, a fitness is calculated (line 14) as,

$$F_M = \bar{p} \cdot tp \cdot (1 + w_A \cdot \sum_{\pi \in \mathcal{P}_v} (a_\pi^U - a_\pi^{max}) / (a_\pi^{max} \cdot 0.01)) \cdot (w_R \cdot \prod_{T \in \Theta_v} t_T / t_T^{max}) \quad (5.4)$$

where the average power dissipation \bar{p} is given by Equation (5.1) and tp introduces a timing penalty if the schedule exceeds task deadlines or the repetition period. Further, an area penalty is applied for all PEs with area violation \mathcal{P}_v by relating used area a_π^U and area constraint a_π^{max} . Similarly, a transition time penalty is applied for all transitions Θ_v that exceed their maximal transition time limit, i.e., transition time t_T exceeds the maximal allowed transition time t_T^{max} . Both area and transition penalty are weighted (w_A and w_R), which allows to adjust the aggressiveness of the penalty. Having assigned a fitness to all individuals of the population, they are ranked using linear scaling (line 15). A tournament selection scheme is used to pick individuals (line 16) for mating (line 17). The produced offsprings are inserted into the population (line 18). In order to push the GA away from infeasible and low quality design space regions, four new genetic mutation operators are applied (lines 19–22). These mutation strategies are introduced next.

Shutdown Improvement: To increase the chances of component shutdown, which leads to a reduction of static power consumption, the genetic task mapping algorithm

employs a simple yet effective strategy during the optimisation. Out of the current population randomly picked individuals (probability 2% was found to lead to good results) are modified as follows. A single mode O_x and a non-essential PE π_a are selected. Non-essential PEs are considered to be PEs that implement task types that have alternative implementations on other PEs, hence, they are not fundamental for a feasible solution. Our goal is to switch off PE π_a during the execution of mode O_x . Therefore, all tasks of mode O_x which are mapped to π_a are randomly re-mapped to the remaining PEs ($\mathcal{P} \setminus \pi_a$), hence, PE π_a can be shut down during mode O_x . Of course, only feasible mappings are allowed, i.e., tasks are always mapped randomly to the PEs that are capable of executing this kind of task type. The pseudo code of this shutdown improvement strategy is shown in Figure 5.8. A list of non-essential PE, i.e., PEs that have the potential to be switched off, is produced in line 01. The operational mode O_x is selected randomly, however, the greater the execution probability of a mode the higher are the chances to select this mode (line 02–08). The PE to be switched-off π_a is randomly selected in the lines 09–12. PEs accommodating less tasks than others are selected with a higher probability, since re-mapping only a small number of tasks is likely to influence the execution behaviour less drastically than the re-mapping of many tasks. After the operational mode O_x and the PE π_a have been selected, all tasks mapped to PE π_a in mode O_x are randomly re-mapped to the remaining PEs (line 13–14). The so mutated string is returned and inserted into the population pool. Due to the functional similarity of the four improvement strategies, only the pseudo code of the shutdown technique is given here (see Figure 5.8).

Area Improvement: To avoid convergence towards area infeasible solutions, a second strategy is employed. If only infeasible area mappings have been produced for a certain number of generations, the search is pushed away from this region by randomly re-mapping hardware tasks onto software-programmable PEs.

Timing Improvement: In contrast to the area improvement strategy, if a certain amount of timing infeasible solutions have been produced, software tasks are randomly mapped to faster hardware implementations. Thereby, the chance to find timing feasible implementations is increased.

Transition Improvement: Cores implemented in FPGAs can be dynamically reconfig-

Algorithm: SHUTDOWN-IMPROVEMENT

Input: - Randomly picked initial mapping string M (probability 2%)
 - OMSM (finite state machine + task graphs)
 - Technology Library, Allocated Architecture

Output: - For shutdown modified mapping string

```

01: NonEssPEs  $\leftarrow$  PossibleShutdownPEs()
02: Rand = RandomFloat(0, 1)
03: AccuProb = 0
04:  $O_x = 0$ 
05: forall (Mode, ModeNumbers) {
06:   AccuProb += GetExecutionProb(Mode)
07:   if (RandDouble  $\leq$  AccuProb) break
08:   else  $O_x++$  }
09: Rand = RandomFloat(0, TasksNoInvSum)
10: forall(NoOfPE, NoAllocPEs) {
11:   if ((InvInterval0[NoOfPE]  $\leq$  RandDouble) &&
        (InvInterval1[NoOfPE]  $>$  RandDouble)) {
12:      $\pi_a = \text{NoOfPE}$  } }
13: forall (gene  $\neq \pi_a$ )  $\in M$ 
14:   gene = RandomFeasibleMapping( $\mathcal{P} \setminus \pi_a$ )
15: return  $M$  // modified mapping

```

Figure 5.8: Pseudo Code: Mapping Modification towards component shutdown

ured. However, this involves a time overhead. If this overhead exceeds the imposed transition time limits, the mapping is infeasible. Hence, after generating for a certain number of generations solely solutions that violate the transition times, tasks are randomly re-mapped away from the FPGAs that cause the violations.

Although some of the produced genomes (strings) might be infeasible in terms of area and timing behaviour, all these strategies have been found to improve the search process significantly by introducing individuals that evolve into high quality solutions. For instance, running the synthesis process (on examples of moderate size) without the shutdown improvement strategy often results in implementations which do not exploit this energy reduction possibility.

5.5.2 Hardware Core Allocation

For tasks mapped to ASICs and FPGAs it is necessary to allocate hardware cores that are capable of executing the task types. This is a trivial job as long as only tasks of different types are mapped to the same hardware component, i.e., when a single core for each task needs to be allocated. Nevertheless, if tasks of the *same* type η are assigned to the *same* PE more than once, it is necessary to make a decision upon how many core of type η need to be implemented. This is important because hardware cores are able to execute tasks in parallel, i.e., the right quantitative choice of cores can efficiently help to exploit application parallelism, hence, improve the timing behaviour as well as energy dissipation. In the proposed co-synthesis, the following approach is employed during the schedule optimisation. Initially, each task type assigned to hardware is implemented only once, even if multiple tasks of this type are mapped onto the same PE. This ensures that all hardware tasks have at least one executable core implementation. If the hardware area constraints are not violated through the initial allocation, additional cores are implemented as follows. The tasks are analysed to identify possibly parallel executing task, taking into account task dependencies. These tasks are then ordered according to their mobility. Clearly, tasks with low mobility are more likely to improve the timing behaviour and therefore should be the preferred choice when implementing additional hardware cores. Accordingly, cores for tasks with low mobility are implemented as long as the area constraints of the hardware components are not violated. Note, this strategy potentially improves the energy dissipation, since it is probable to result in more slack time, which, in turn, can be exploited through DVS.

5.5.3 Dynamic Voltage Scaling for Multiple Parallel Executing Tasks

Dynamic voltage scaling is a powerful technique to reduce energy consumption by exploiting temporal performance requirements through dynamically adapting processing speed and supply voltage of PEs. An effective voltage scaling technique for this purpose has already been introduced in Chapter 3. Furthermore, the applicability of DVS to embedded distributed systems was demonstrated in [64, 90]. However, these works as well as the energy-gradient technique introduced in Chapter 3 concentrate on dynamically changing the performance of software PEs only, while parallel execution of tasks on hardware resources has been neglected. Nevertheless, in the context of energy-efficient




Figure 5.9: DVS Transformation for HW Cores

multi-mode systems, where performance requirements of each operational mode can vary significantly, DVS needs to be considered carefully. Consider, for instance, an inverse discrete cosine transformation (IDCT) algorithm implemented in fast hardware which is used during two modes: MP3 decoding and JPEG image decoding. Clearly, the JPEG decoder should restore images as quickly as possible, i.e., the IDCT hardware is required to execute at maximal supply voltage (equivalent to peak performance). On the other hand, the MP3 decoder works at a fixed repetition rate of $25ms$ for which the hardware implementation operates faster than necessary, i.e., the IDCT performance can be reduced such that this repetition rate is adequately met. By using DVS it is possible to adapt the execution speed to suit both needs and to reduce the energy consumption to a minimum.

This chapter considers that hardware components might employ DVS. However, due to the area and power overhead involved in additional DVS circuitry (DC/DC-converter [61, 97]) it is assumed that all cores allocated to the same hardware component are fed by a single voltage supply, i.e., dynamically scaling the supply voltage simultaneously affects the performance of all cores on that hardware component. To cope with this problem, the potentially parallel executing tasks on a single scalable hardware resource are transformed into an equivalent set of sequentially executing tasks, taking into account the dynamic power dissipation on each core. Note that this is done to calculate the scaled supply voltages only, i.e., this virtual transformation does not affect the real implementation. Figure 5.9 shows the transformation of five hardware tasks, executing on two cores (both cores are implemented within the same hardware component), to three sequential tasks on a single core. This sequential execution is equivalent to the behaviour of software tasks, hence, a voltage scaling technique for software processors can be applied. Nevertheless, it

is important to consider task dependencies and task deadlines during this transformation, i.e., it might be necessary to further subdivide the transformed task graph, in order to maintain a correct specification. For instance, consider original specification given in Figure 5.10. Here, task τ_3 has a data dependency with another task outside of the given

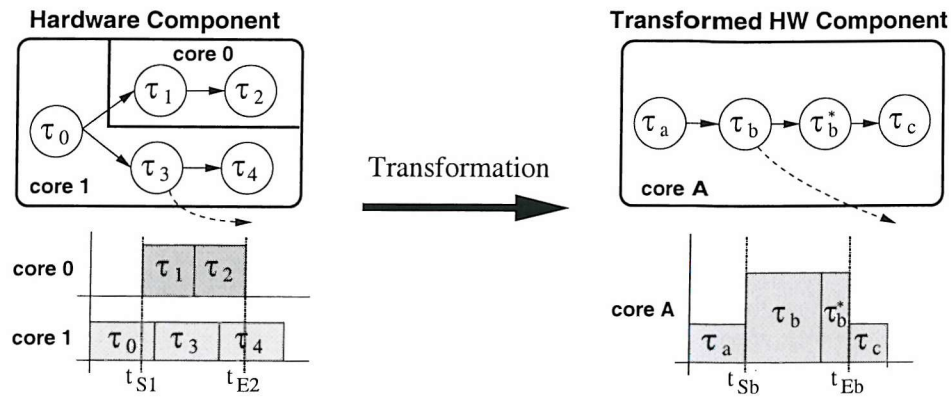


Figure 5.10: DVS Transformation for HW Cores considering inter-PE communication

component. Due to the data dependencies the transformation results in four tasks, instead of three (Figure 5.9). Thereby, task τ_b and task τ_b^* can be scaled independent of each other. While task τ_b can influence tasks on other PEs, task τ_b^* solely affects downstream tasks. Exactly as the original specification. Task deadlines can be handled similarly.

The pseudo code description of this transformation is given in Figure 5.11. The function takes as input the system task graph and a hardware processing element that is DVS enabled. In the first two steps (lines 1 and 2) it generates two priority queues of tasks mapped to the HW processing element, sorted in decreasing order of start and end time, respectively. A sequential task graph and a current power variable are initialised (lines 3 and 4). At the end of algorithm run, the sequential task graph will hold the transformed tasks. The following transformations are repeated until no tasks are left in the priority queue of task sorted according to their end times (`EndTaskQueue`):

- If the next chronological event is due to a starting task (line 6), the current power dissipation is increased by the power dissipation of the starting task or tasks (line 7) and a new task is added to the sequential task graph by inheriting the current power dissipation and the start time (line 8). The starting tasks are removed from `StartTaskQueue` (line 9).
- If the next chronological event happens due to an end task (line 11), then the current

Algorithm: DVS-HARDWARE-TRANSFORMATION

Input: - Task Graphs
 - DVS-enabled Hardware Processing Element
Output: - Sequentialised tasks on the HW PE

```

01: StartTaskQueue = InitStartTaskQueue(Tasks)
02: EndTaskQueue = InitEndTaskQueue(Tasks)
03: SeqTaskGraph  $\leftarrow$  0
04: Pw  $\leftarrow$  0
05: while (EndTaskQueue  $\neq$  0) {
06:   if (TimeStartTask < TimeEndTask) {
07:     Pw += Pw(StartTasks)
08:     AddTask(SeqTaskGraph, Pw, StartTime)
09:     RemoveStartTasks(StartTaskQueue)
10:   }
11:   else if (TimeEndTask < TimeStartTask) {
12:     Pw -= Pw(EndTasks)
13:     if (Pw  $\neq$  0)
14:       AddTask(SeqTaskGraph, Pw, OutEdgesOffPE)
15:     else
16:       InheritEndTime(SeqTaskGraph->Last, OutEdgesOffPE)
17:     RemoveEndTasks(EndTaskQueue)
18:   }
19:   else if (TimeStartTask = TimeEndTask) {
20:     if (Core(StartTasks) = Core(EndTasks)) {
21:       if (OutEdgesOffPE(EndTasks)  $\neq$  0)
22:         AddTask(SeqTaskGraph, Pw, OutEdgesOffPE)
23:     }
24:     else { // Core(StartTasks)  $\neq$  Core(EndTasks) {
25:       Pw = Pw + Pw(StartTasks) - Pw(EndTask)
26:       AddTask(SeqTaskGraph, Pw, OutEdgesOffPE)
27:     }
28:     RemoveStartTasks(StartTaskQueue)
29:     RemoveEndTasks(EndTaskQueue)
30:   }
31: }
32: return (SeqTaskGraph)

```

Figure 5.11: Pseudo code: Task graph transformation for DVS-enabled hardware cores

power dissipation is decreased by the power of the ending tasks (line 12). According to whether the current power has reach zero or not, a task is added to the sequential task graph by inheriting the current power value and the outedges (if any) of the ending task (line 14), or the end time and the outedges are inherited to the last

task in the sequential task graph (line 16). The ending tasks are removed from `EndTaskQueue` (line 17).

- If the next chronological events occur to the simultaneously ending and starting task the following procedure is necessary (line 16). If both events appear on the same HW core (i.e., they subsequent tasks), and there are outedges emitting from the ending task, then a new task is added to the sequential task graph. However, the current power is maintained at the same level. Nevertheless, if the two events occur on different cores the current power is adjusted and an new task is introduced to the sequential task graph. Starting and ending tasks are removed `StartTaskQueue` and `EndTaskQueue` (lines 28 and 29).

After no tasks are left in the task queue, which is sorted in chronologically decreasing order of end times, the sequential task graph is returned (line 32).

5.6 Experimental Results: Multi-Mode

Based on techniques and algorithms presented in this chapter, the multi-mode synthesis approach has been implemented on a Pentium III/1.2GHz Linux PC. In order to evaluate its capability to produce high quality solutions in terms of energy consumption, timing behaviour, and hardware area requirements, a set of experiments has been carried out on 15 automatically generated multi-mode examples (`mul1–mul15`¹) and one real-life benchmark example (`smart-phone`). The following two sections split these experiments into hypothetical and real-life examples. All reported results were obtained by running the optimisation processes 40 times and averaging the outcomes.

5.6.1 Hypothetical Examples

Each of the 15 generated examples (`mul1–mul15`) is specified by 3 to 5 operational modes, each consisting of 8 to 32 tasks. The used target architectures contain 2 to 4 heterogeneous PEs, some of which are DVS enabled. These PEs are interconnected through 1 to 3 communication links. The execution probabilities of individual modes were randomly

¹These examples were generated with the publicly available tool TGFF [44].

Example (No. of modes)	Mode Execution Probabilities	w/o probabilities		with probabilities		
		Average power (mW)	CPU time (s)	Average power (mW)	CPU time (s)	Reduction (%)
mul1 (4)	5:10:75:10	8.131	20.7	7.529	24.7	7.29
mul2 (4)	10:5:80:5	3.404	15.5	2.771	18.2	18.61
mul3 (5)	7:3:80:5:5	10.923	23.4	10.430	23.0	4.17
mul4 (5)	1:4:5:40:50	7.975	21.0	6.726	25.2	15.50
mul5 (4)	7:13:35:45	5.186	18.4	4.668	22.1	10.01
mul6 (4)	15:10:10:65	1.677	20.6	1.301	19.9	22.46
mul7 (4)	5:5:5:85	3.306	11.6	1.250	21.4	62.18
mul8 (4)	75:5:15:5	1.565	32.1	1.329	28.0	15.06
mul9 (4)	7:3:80:10	3.081	6.0	1.901	5.8	38.28
mul10 (5)	45:5:40:5:5	1.105	28.3	0.941	32.1	14.83
mul11 (3)	80:10:10	2.199	9.3	1.304	16.6	40.70
mul12 (4)	15:10:50:25	7.006	25.4	5.975	34.2	14.69
mul13 (3)	5:15:80	4.090	15.8	2.816	15.8	31.04
mul14 (5)	5:5:10:10:70	8.195	28.6	6.466	33.0	21.13
mul15 (5)	5:10:75:7:3	2.188	41.5	1.222	55.4	44.16

Table 5.2: Considering mode execution probabilities (excluding DVS)

chosen and varies between 1% and 85%. To illustrate the importance of taking mode execution probabilities into account during the synthesis process, an execution probability neglecting approach is compared with the proposed synthesis technique, which considers the mode probabilities. The first two sets of experiments demonstrate the energy savings achievable through the consideration of mode execution probabilities, either with or without the exploration of DVS. While the third set examines the influence of activation profile on the energy savings.

Comparisons excluding Dynamic Voltage Scaling

To highlight the influence of mode execution probabilities on the achievable energy saving consider Table 5.2 which shows the multi-mode co-synthesis results for the 15 automatically generated benchmarks. The first two columns give the benchmark names and the mode execution probabilities. The third and the fourth column present the dissipated average power and optimisation time for the execution probability neglecting synthesis approach, while the fifth and sixth column show the same for the proposed approach. Take, for instance, example mul6. When ignoring the execution probabilities during the

optimisation, an average power dissipation of $1.677mW$ is achieved. However, optimising the same benchmark example under the consideration that modes execute with *uneven* probabilities (e.g., 15:10:10:65 — i.e., mode 1 is active for 15%, mode 2 is active of 10%, and so on), the average power can be reduced by an appropriate task mapping and core allocation to $1.301mW$. This is a significant reduction of 22.46%. Furthermore, it can be observed that the proposed technique was able to reduce the energy consumption of all examples with up to 62.18% (mul7). Note that these reductions are achieved without any modification of the underlying hardware architectures, i.e., the system costs are not increased. It is also important to note that the achieved energy reductions are solely introduced by taking the mode execution probabilities into account during the co-synthesis process, i.e., both compared approaches allow the same resource sharing and rely on the same scheduling technique. When comparing the optimisation times for both approaches, it can be observed that the proposed technique shows a slightly increased CPU time for most examples, which is mainly due to the design space structure.

Comparisons including Dynamic Voltage Scaling

The next experiments were conducted to see how the proposed technique compares to DVS and if further savings can be achieved by taking the mode probabilities and DVS simultaneously into account. Table 5.3 reports on the findings.

The DVS technique that was used here is based on PV-DVS (Chapter 3), which has been extended to enable the consideration of DVS not only for software processors, but also for parallel executing cores on hardware PEs (see Section 5.5.3). As in the first experiments, two approaches are compared here. The first approach disregards the mode execution probabilities during optimisation, while the second takes them into account throughout the co-synthesis. Similar to Table 5.2, the second and the third column of Table 5.3 show the results without consideration of execution probabilities, whilst the fourth and the fifth column present the results achieved by taking execution probabilities into account. Lets consider again benchmark mul6. Although the execution probabilities are neglected in the second column, a reduced average power consumption ($0.689mW$) can be observed, when compared to the results given in Table 5.2. This clearly demonstrates the high energy reduction capabilities of DVS. Nevertheless, it is possible to further minimise the power consumption to $0.465mW$ by considering the execution probabilities together with DVS. This is an improvement of 32.53%, solely due to the synthesis for the partic-

Example (No. of modes)	Mode Execution Probabilities	w/o probabilities		with probabilities		
		Average power (mW)	CPU time (s)	Average power (mW)	CPU time (s)	Reduction (%)
mul1 (4)	5:10:75:10	4.271	526.6	3.964	768.6	10.92
mul2 (4)	10:5:80:5	1.568	860.4	1.273	687.4	18.82
mul3 (5)	7:3:80:5:5	4.012	1053.5	3.344	1192.2	16.66
mul4 (5)	1:4:5:40:50	2.914	1135.2	2.320	1125.4	20.39
mul5 (3)	7:13:35:45	1.394	967.7	1.315	932.1	5.68
mul6 (4)	15:10:10:65	0.689	472.9	0.465	593.7	32.53
mul7 (4)	5:5:5:85	1.331	540.3	0.479	820.7	64.02
mul8 (4)	75:5:15:5	0.564	1262.1	0.436	1412.0	22.64
mul9 (4)	7:3:80:10	0.942	161.2	0.648	177.1	34.66
mul10 (5)	45:5:40:5:5	0.480	1456.3	0.394	1361.9	17.88
mul11 (3)	80:10:10	0.396	318.1	0.255	403.2	35.53
mul12 (4)	15:10:50:25	2.857	1384.7	2.460	1450.7	13.91
mul13 (3)	5:15:80	1.185	498.3	0.953	576.6	19.56
mul14 (5)	5:5:10:10:70	2.320	1512.3	1.797	1556.4	22.55
mul15 (5)	5:10:75:7:3	0.801	1316.7	0.324	1836.3	59.56

Table 5.3: Considering mode execution probabilities (including DVS)

ular execution probabilities. For all other benchmarks savings of up to 64.02% (mul7) were achieved. Due to the computation of scaled supply voltages and the influence of scheduling on the energy consumption, the optimisation times are higher when DVS is considered.

Figure 5.12 shows the Pareto outcome of the typical co-synthesis run. Each diamond in this graph indicates a possible implementation of benchmark mul15, however, each with a different trade-off between energy dissipation of area usage. The rightmost solution respects all imposed system constraints including area (area violation is 0). This solution dissipates an energy of 1.179mW. Nevertheless, the graph indicates that further energy savings can be made by mapping more functions towards hardware. Of course, increasing the used area corresponds to more costly design. For instance take the leftmost solution, which dissipated 0.773mW. To achieve this reduced energy dissipation it is necessary to increase the available hardware area by 124%, i.e., the hardware requirements are more than doubled.




Figure 5.12: Pareto optimal solution space achieved through a single optimisation run of `mul15` (without DVS), revealing the solution trade-offs between energy dissipation and area usage

Influence of Real Activation Probabilities

The next experiment is conducted to highlight the influence of the user behaviour on the energy efficiency of a system that has been synthesised under the consideration of certain mode execution probabilities. Certainly, the mode execution probabilities, which are used during the synthesis represent an "imaginative" user and the activation probabilities of a *real* user will differ from those. Accordingly, the following experiment tries to answer the question how the energy efficiency is affected by different activation profiles during application run-time. For experimental purpose a simple specification with two modes is used which contains 14 and 24 tasks (mode 1 and mode 2). The underlying architecture consists of two programmable processors and a single ASIC, all connected via a shared bus. This configuration was synthesised for three different pairs of execution probabilities (0.1:0.9, 0.9:0.1, and 0.5:0.5). These three implementations possibilities correspond to the three lines shown in Figure 5.13. All implementations are based on the same hardware architecture; yet, each has a different task and communication mapping, core allocation, as well as schedule. The first solution (red) was synthesised under the consideration of execution probabilities 0.1:0.9, that is, it is assumed that mode 1 and mode 2 are active




Figure 5.13: A system specification consisting of two operational modes optimised for different execution probabilities (red–0.1:0.9, blue–0.9:0.1, green–0.5:0.5)

for 10% and 90% of the operational time, respectively. Similarly, the second (blue) and the third (green) line represent solutions that have been synthesised using execution probabilities 0.9:0.1 and 0.5:0.5, respectively. According to the real execution probabilities during run-time, i.e., the activation behaviour of the user, the average power dissipations of the implemented systems vary. Consider the system optimised for execution probabilities 0.1:0.9 (red line). If the user behaviour corresponds to these probabilities (User A), the system dissipates an average power of approximately 5.5mW (point I). However, if a different user (User B), for instance, uses mode 1 for 90% and mode 2 for 10% of the time (0.9:0.1), the system will dissipate approximately 12.5mW (point II). Nevertheless, if the system would be optimised for this activation profile (0.9:0.1), as indicated by the blue line in Figure 5.13, a lower power dissipation of around 2.6mW (point III) can be achieved. Similarly, if the system is optimised for execution probabilities 0.9:0.1 (blue line) and the user runs the application 10% in mode 1 and 90% in mode 2 (User A), then a power dissipation of 12.5mW (point IV) is given. While an optimisation towards this usage profile can achieve a system implementation which dissipates only 5.5mW (point I), i.e., extending the battery-lifetime by a factor of 2.83 times. The green plot in Figure 5.13

represents the solution when the execution probabilities are neglected during the optimisation, that is, the execution probabilities are considered to be equal for both modes. Of course, if the modes 1 and 2 are active for equal amounts of time, this solution achieves a lower power dissipation ($6.5mW$, point V) than the systems optimised for execution probabilities 0.1:0.9 ($9mW$, point VI) and 0.9:0.1 ($7.6mW$, point VII). In summary, Figure 5.13 clearly shows that the execution probabilities substantially influence the energy dissipation of the system. Certainly, the system should be optimised as close as possible towards the real behaviour to achieve low energy consumptions, which, in turn, result in longer battery-lifetimes.

5.6.2 Smart Phone Benchmark

To further validate the co-synthesis technique in terms of real-world applicability, the introduced approach was applied to a smart-phone example. This benchmark is based on three publicly available applications: a GSM codec [4], a JPEG codec [5], and an MP3 decoder [68]. Accordingly, the smart-phone offers three different services to the user, namely, a GSM cellular phone, a digital camera, and an MP3-player. Of course, the used applications do not specify the whole smart-phone device, however, a major digital part of it. The specification for this example, given as operational mode state machine (OMSM), has already been introduced in Figure 5.1 (page 105). For each of the eight operational modes, the corresponding task graphs have been extracted from the above given references. Details about these task graphs specifications can be found in Appendix B. The individual applications have been software profiled to gather the necessary execution characteristics of each task. This was carried out by compiling profile information into the application [1, 3] and running the produced software on real-life input streams. On the other hand, the hardware estimations are not based on direct measurements, but have been based on typical values, such that hardware tasks typically executed 1 to 2 orders of magnitude faster and dissipated 1 to 2 orders of magnitude less power than their software counterparts [32]. Depending on the operational mode, the number of tasks and communications vary between 5–88 nodes and 0–137 edges, respectively. The hardware architecture of the embedded system within the smart phone consists of one DVS-enabled processor (based on the ARM8 developed in [31]) and two ASICs. These components are connected via a single bus. Table 5.4 gives the results of the conducted experiments, which are also graphically shown in Figure 5.14. Similar to the previous experiments,

	without probabilities		with probabilities		
	Average power (<i>mW</i>)	CPU time (<i>s</i>)	Average power (<i>mW</i>)	CPU time (<i>s</i>)	Reduction (%)
w/o DVS	2.602	80.1	1.801	96.9	30.76
with DVS	1.217	3754.5	0.859	4344.8	29.41

Table 5.4: Results of smart phone experiments

an approach which neglects the execution probabilities is compared with the introduced co-synthesis technique that considers the uneven activation times of different modes. The first row in Table 5.4 shows this comparison for a fixed voltage system, i.e., no DVS is applied. Synthesising the system without consideration of execution probabilities results in an average power consumption of $2.602mW$, when running the system after the synthesis according to the activation profile. Nevertheless, taking into account the mode usage profile during the co-synthesis this can be reduced by 30.76% to $1.801mW$. Again, this saving is achieved without the modification of the allocated hardware architecture, therefore, the system cost is the same for both solutions. Also DVS has been applied to this benchmark, considering that the GPP of the given architecture supports DVS functionality. The results are shown in the second row of Table 5.4). It can be observed that the power consumption of the smart phone drops to $1.217mW$, even when neglecting mode execution probabilities. However, the combination of applying DVS and taking execution probabilities into account results in the lowest power consumption of $0.859mW$, a 29.41% reduction, when compared to the DVS neglecting approach. Overall, the average power is decreased from $2.602mW$ to $0.859mW$. Which represents a significant reduction of nearly 67%. Regarding the required co-synthesis times, the four implementations could be found in 80.1s (without probabilities and DVS) to 4344.8s (with probabilities and DVS). Clearly, considering DVS requires longer optimisation times due to the voltage scaling problem that needs to be solved repetitively within the innermost optimisation loop of the co-synthesis algorithm. For instance, the optimisation for DVS increase the run-time from 80.1s to 3754.1s for the case without consideration of execution probabilities, and from 96.9s to 4344.8s when execution probabilities are taken into account. On the other hand, the consideration of mode execution probabilities increases the optimisation time only moderately form 80.1s to 96.9s in the case of no DVS, and from 3754.5s to 4344.8s if the probabilities are considered.




Figure 5.14: Energy dissipation of the Smart phone using different optimisation strategies

5.7 Concluding Remarks

This chapter has introduced new techniques and algorithms for the energy minimisation of multi-mode embedded systems. An abstract specification model called operational mode state machine has been proposed. This model allows for the concurrent specification of mode interaction (top-level finite state machine) as well as mode functionality (task graph). The advantage of such a representation is the capability to express the complete functionality of the system within a single model, containing both control and data flow.

The presented co-synthesis technique not only optimises mapping and scheduling towards hardware cost and timing behaviour, but also aims at the reduction of power consumption at the same time. A key contribution in this chapter has been the development of an effective mapping strategy that considers uneven mode execution probabilities as well as important power reduction aspects, such as multiple task implementations and core allocation. For this purpose a GA based mapping approach has been proposed along with four improvement strategies to effectively handle the optimisation of component shutdown, transition time, area usage, and timing behaviour. These improvement strategies guide the mapping optimisation of multi-mode specifications towards high quality solutions in terms of power consumption, timing feasibility, and area usage. Furthermore, these strategies help to reduce the optimisation times needed by conventional genetic algorithms, since they use computational inexpensive constructive heuristic in a local search

fashion. A newly introduced transformational-based algorithm for DVS-enabled hardware components, which is capable of performing parallel task execution, allows to easily leverage the efficiency of existing voltage scaling algorithms. This algorithm transforms a set of potentially parallel executing tasks on a single HW component into a set of sequential executing tasks, taking into account imposed deadlines and inter-PE communications.

The proposed techniques and algorithms have been validated through extensive experiments including a smart phone real-life example. These experiments have demonstrated that taking into account mode execution probabilities throughout the system synthesis leads to substantial energy savings compared to conventional approaches which neglect this issue. Furthermore, DVS has been considered in the context of multi-mode embedded systems and it was shown that considerably high energy reductions can be achieved by combining both the consideration of execution probabilities and dynamic voltage scaling.

Chapter 6

LOPOCOS: A Prototype Low Power Co-Synthesis Tool

Chapters 3, 4, and 5 have described a number of new techniques for the design of energy-efficient distributed embedded systems. The main aim of this chapter is to show how these techniques can be used to explore the architectural design space with the intention of finding high quality system-level designs. The intention of the architectural design space exploration is the identification of a suitable architecture that will be integrated inside a new product, as outlined in Chapter 1. For this purpose the prototype co-synthesis tool, LOPOCOS (Low Power Co-Synthesis), has been developed. LOPOCOS incorporates the algorithms proposed in Chapters 3, 4, and 5. The overall goal of LOPOCOS is to equip system designers with a tool that helps to effectively explore different architectural implementations through an automated system-level design process. Using LOPOCOS, it will be demonstrated how suitable architectures (combinations of processing elements which are interconnected via communications links) can be derived for a realistic smart phone example. The smart phone example has been chosen because it is of sufficient complexity to assess the capability of LOPOCOS in terms of real-world applicability.

The remainder of this chapter is organised as follows. Section 6.1 discusses the smart phone applications and their representations as task graphs. Section 6.2 introduces LOPOCOS and demonstrates its usage for architectural design space exploration. Finally, Section 6.3 provides some concluding remarks.

6.1 Smart Phone Description

Emerging smart phones often combine cellular phones, digital cameras, and MP3 players into a single device. The specification through task graphs of such a smart phone is described in this section. The section is split into three subsections which briefly outline the applications used within the smart phone. Section 6.1.1 describes the voice compression and decompression algorithms used within GSM (Global System for Mobile Telecommunication) cellular phones. Section 6.1.2 details the decompression algorithm used for the MP3 audio files. And finally, Section 6.1.3 provides information about the JPEG image compression and decompression standard. These applications and their combinations are used to model the different modes in the operational mode state machine of the smart phone (Figure 5.1, Page 105).

6.1.1 Voice Compression and Decompression

For reasonably good communications, the human voice has to be sampled at $8kHz$ and quantised with at least $13bit$. Thus, the transmission of one second of human voice requires a data transfer of $13000bytes$. To reduce this amount of data, GSM cellular phones compress voice streams using regular pulse excitation long-term predictor full-rate speech transcoders, or short RPE-LTP transcoders. These transcoders model the human voice system through two filters, the linear-predictive short-term filter and the long-term predictive filter. Figure 6.1 shows block diagrams of the encoder and the decoder units. The encoder divides the incoming voice signal into short-term predictable parts, long-term predictable parts, and the remaining residual pulse. The resulting parameters of the filters and residual pulse are quantised and encoded. Using this compression technique, the $13000byte/s$ data stream is reduced to $1625byte/s$, a compression ratio of 8 to 1. Upon receiving the encoded voice stream, the decoder decompressed the parameter settings for the filters and the residual pulse. Using these settings the "original" voice signal is reconstructed. A publicly available C implementation of the GSM transcoder was developed by Degner *et al.* [4]. From this description the task graphs for the encoder and the decoder were derived, in order to permit their usage for the co-synthesis. The derivation of the task graph requires careful consideration of several aspects, including:

Loop unrolling: In order to reveal parallelism in the application, complex loops with a

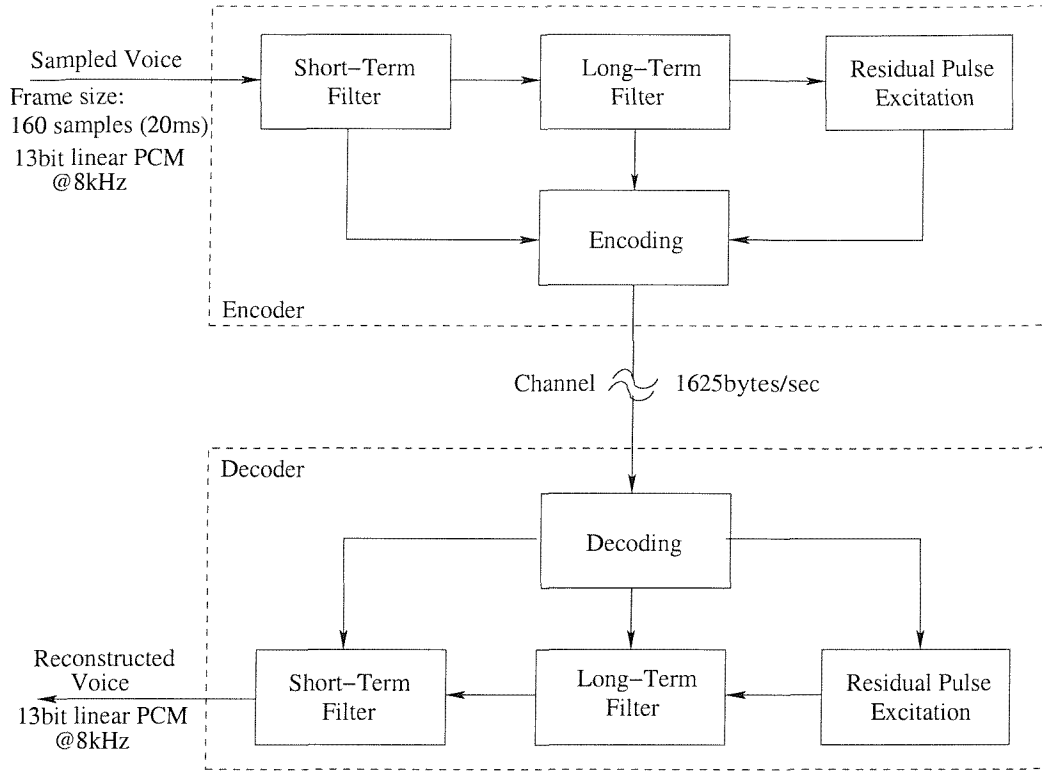


Figure 6.1: Block diagram of the GSM RPE-LTP transcoder [69]

fixed iteration counts can be unrolled and split into several tasks.

Function collapsing: Several trivially small sub-functions are collapsed into single tasks.

Data flow extraction: Memory variables that are shared among functions are transformed into data dependencies between tasks.

Figures 6.2 and 6.3 show the extracted task graph specification of the GSM encoder and decoder. The encoder consists of 53 tasks and 80 edges, while the decoder is given by 34 tasks and 55 data dependencies. In accordance to the input frame size of 160 samples (each covering one sample period of 8kHz), both algorithms need to be performed with a repetition rate of 20ms. If these rate cannot be satisfied, frames are dropped, and each dropped frame corresponds to a voice gap of 20ms, which seriously affects the quality of the communication.

In order to estimate the execution properties of each task, the specifications have been software profiled (assessment of function calls and time spend in each function) using several realistic input voice streams. Details about the execution properties of tasks are given in Appendix B, where the task graphs and the technology files are presented.

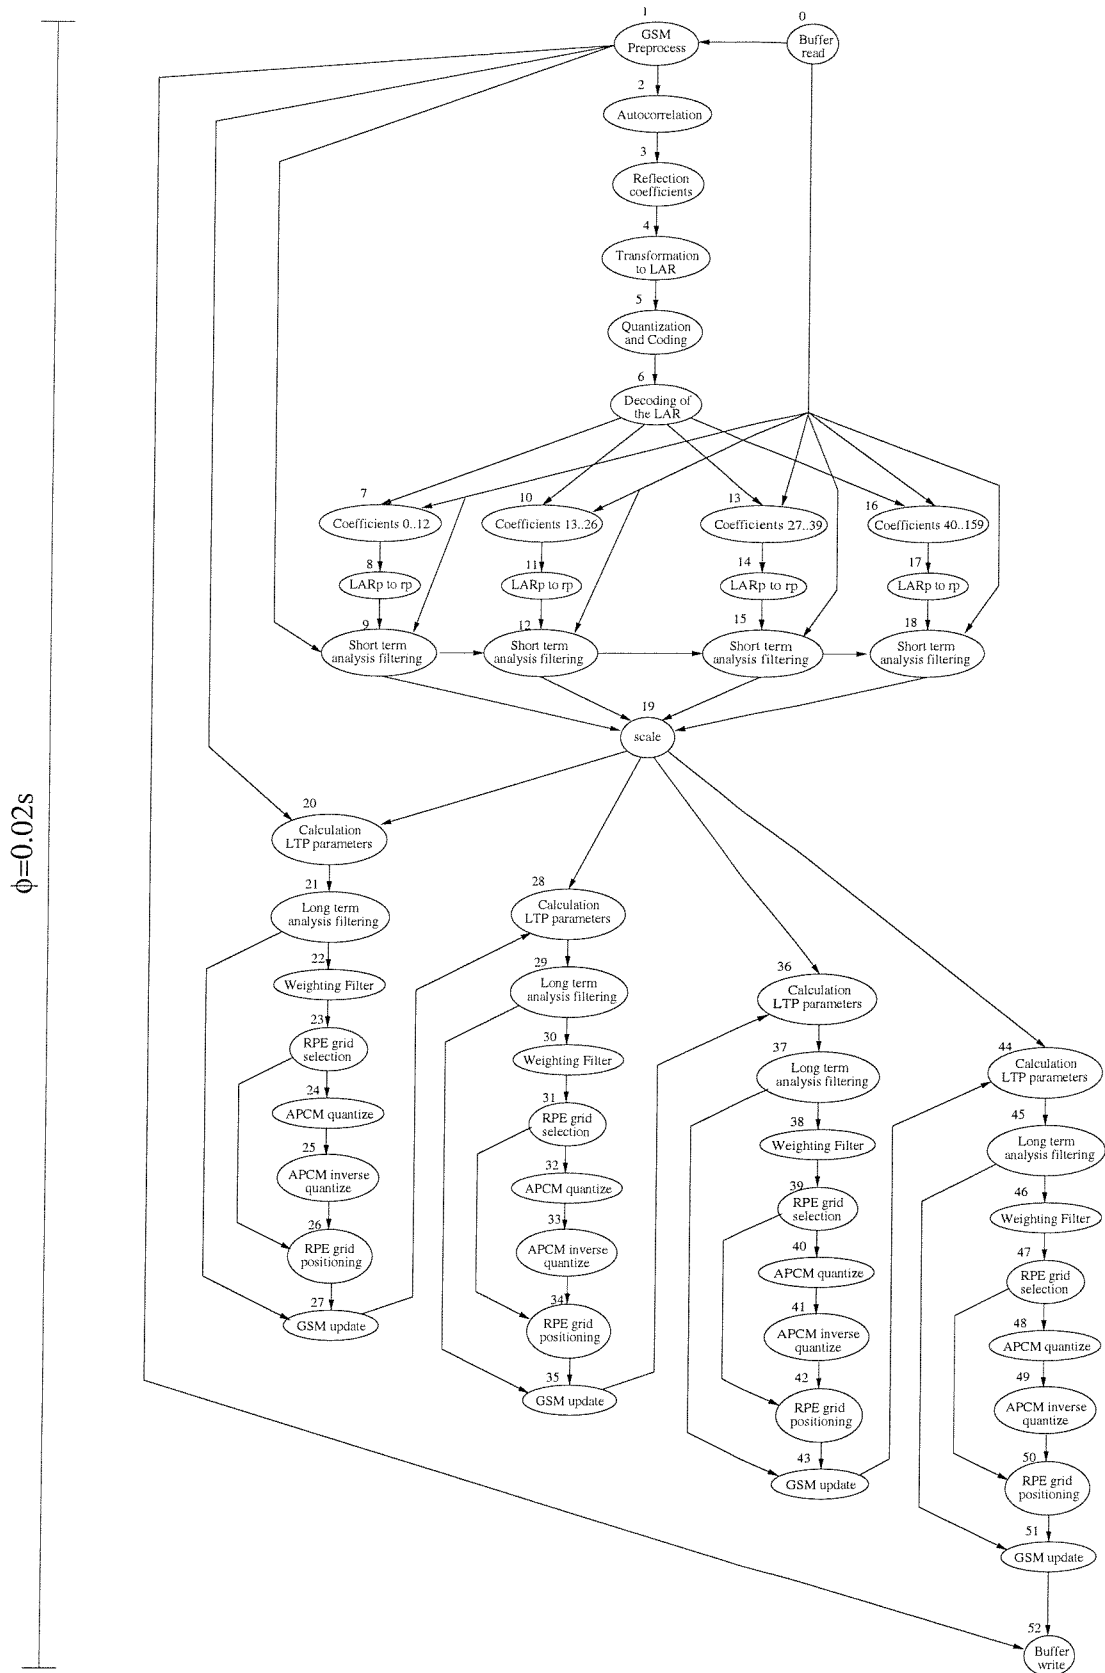


Figure 6.2: Task graph of the GSM voice encoder

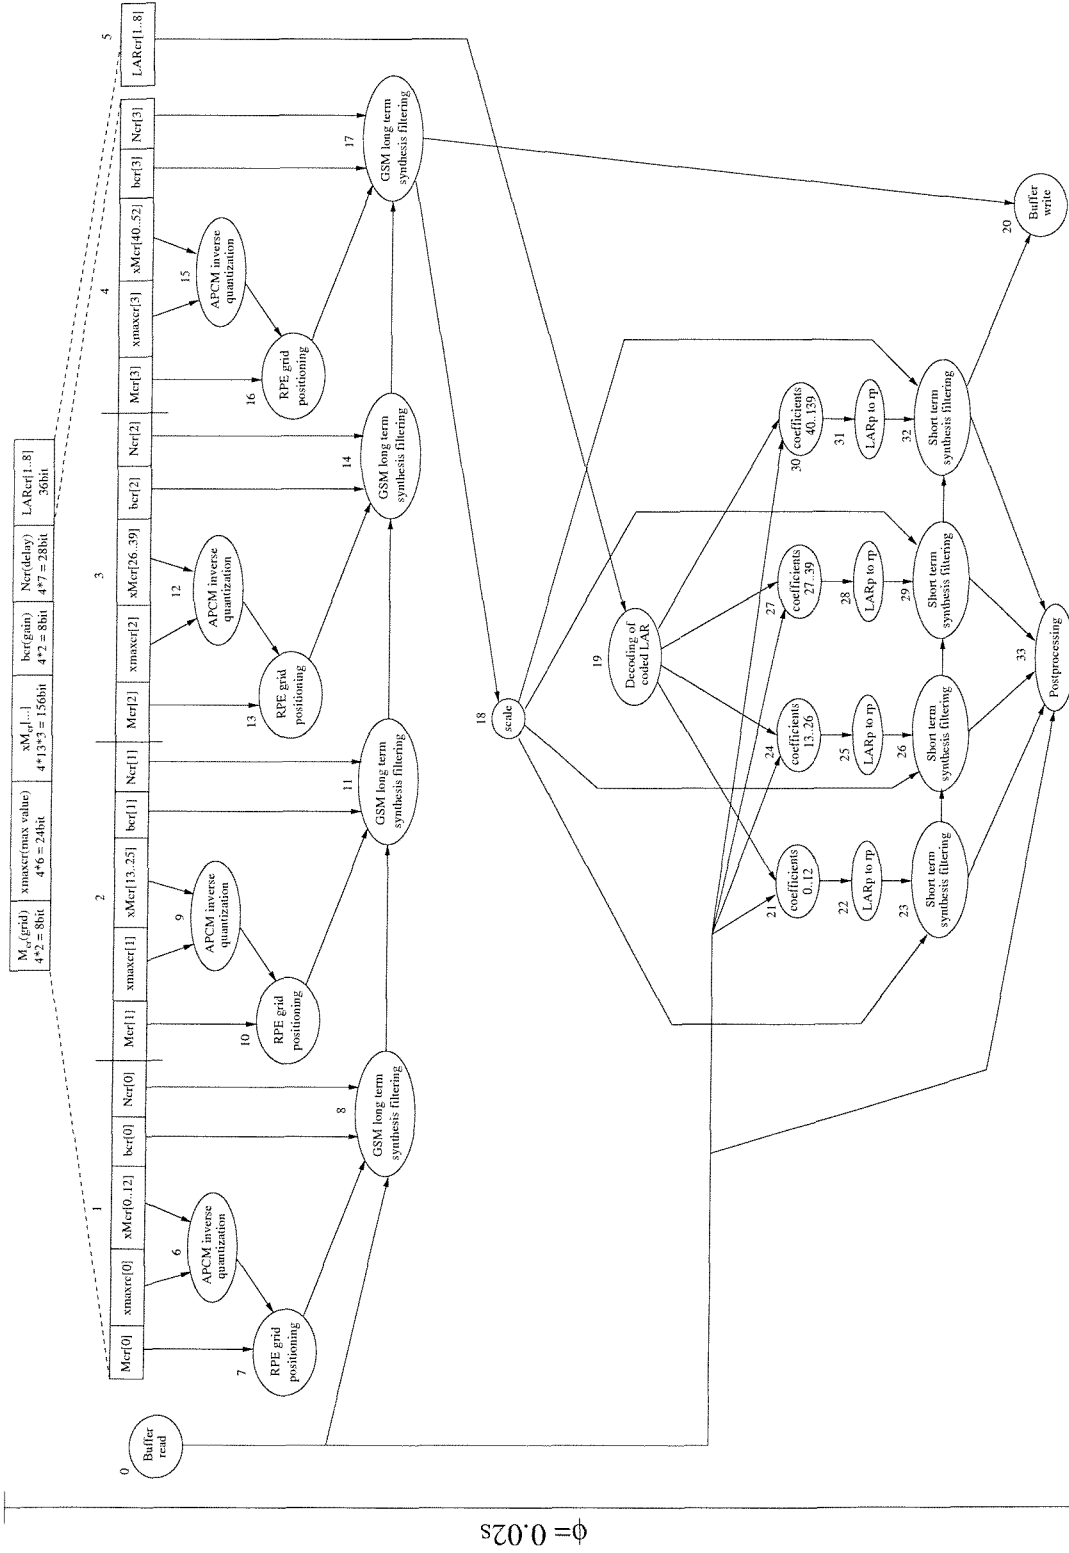


Figure 6.3: Task graph of the GSM voice decoder

6.1.2 MP3 Decoder

MP3 (MPEG-1 Audio Layer 3) has become the *de facto* standard for the compression of music audio signals. Using a bit stream of 128kbps, music can be stored in near-CD quality. The block diagram shown in Figure 6.4 outlines the functionality of the decoder. Each input frame contains a header and the signal samples. The header describes the




Figure 6.4: Block diagram of the MPEG-1 layer 3 audio decoder

encoding parameters such as sample frequency, stereo (mono) signal, and the block type, while the signal samples represent the encoded audio information. In the first step of the MP3 decoder, the input frames are separated into header and samples. The samples are passed to the Huffman decoder and are dequantised to derive the coefficients for the inverse cosine transformation (IDCT). The outcome of the IDCT is fed into the sub-band synthesis filter which reconstructs the "original" audio signal from different frequency sub-bands. Similar to the GSM transcoder, the task graphs specification was derived from a publicly available MP3 decoder implementation [68]. The task graph representation was introduced in Chapter 1 (Figure 1.3, Page 7). It consists of 17 tasks and 18 data dependencies. To ensure an uninterrupted audio signal of high quality at the output, the execution of all tasks in the graph has to be repeated every 25ms. Execution properties were extracted through the profiling of several real music files. More details are given in Appendix B.

6.1.3 JPEG Image Compression and Decompression

The JPEG compression standard is a transformation-based encoding technique that significantly reduces the required storage space of digital images [142]. Inside the smart phone, the JPEG encoder is responsible for the efficient storage of images that are taken with the integrated digital camera. The taken photos do not only need to be stored in compressed form within the smart phone memory, but also need to be restored upon user requests. Block diagrams of both compression and decompression are shown in Figure 6.5. The input to the encoder consists of image sub-blocks with a size of 8x8 pixels. In the first stage

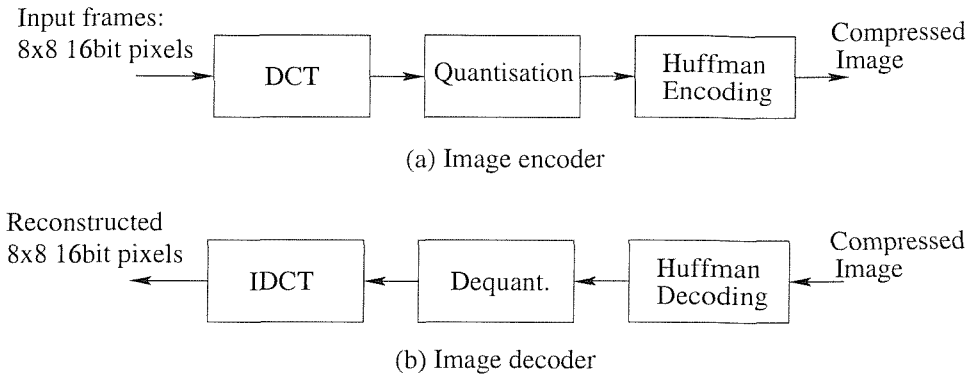


Figure 6.5: Block diagram of the JPEG encoder and decoder [142]

a discrete cosine transformation (DCT) is applied. The goal of the DCT is to redistribute the signal energy to a small set of transformation coefficients. This allows to nullify many of the coefficients during the quantisation, which, in turn, enables an efficient encoding to reduce the storage requirements. The decoder inverts these three steps to reconstruct the image. Due to the quantisation/dequantisation this compression scheme belongs to the class of lossy encoding techniques, i.e., during the encoding some information is "lost" which results in degradation of the image quality. The JPEG library used in smart phone is also publicly available [5]. The corresponding task graphs are given in Figure 6.6, where the sequential nature of the graph can be observed. The encoder needs to store four photos

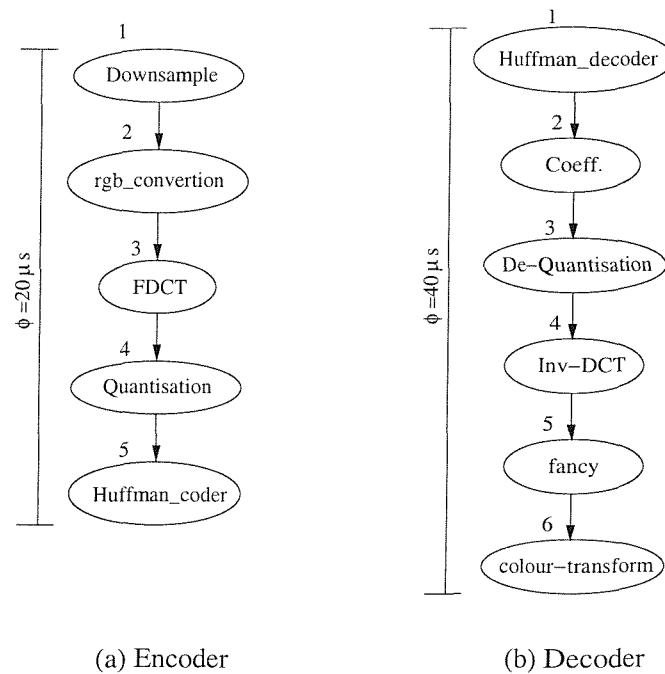


Figure 6.6: Task graphs of the JPEG encoder and decoder

per second to allow a quick series of shots. Hence, the repetition rate should match $20\mu\text{s}$ for the an image size of 1024×768 pixels. The repetition rate of the decoder, on the other hand, needs to be at least $40\mu\text{s}$ in order to restore images within half a second. Similar to the GSM codec and the MP3 decoder, the JPEG encoder and decoder have been profiled on realistic input streams (i.e., photos) to extract task characteristics in terms of execution time. Refer to Appendix B for details about the task properties.

6.2 LOPOCOS

LOPOCOS is a prototype co-synthesis tool that incorporates the algorithms proposed in Chapters 3, 4, and 5. The aim of this section is to show how LOPOCOS can be used for architectural design space exploration. Using the smart phone applications introduced in Section 6.1, the goal of the design space exploration is the identification of an architecture (processing elements connected through communication links) that fulfils the performance requirements and area constraints on one hand, and minimises the energy dissipation and system cost on the other. Figure 6.7 gives an overview of the design flow used within LOPOCOS. The input to LOPOCOS consists of three ASCII files, which contain the following information:

- (a) *System specification* given as task graphs and operational mode state machine (in the case of multi-mode embedded systems).
- (b) *Technology library* containing the available system components that can be used to compose a target architecture. The library describes the execution properties (e.g., execution time, area requirement, dynamic power) of each task when executed on a certain components, as shown Table 6.2. Furthermore, it provides static properties of the available components (e.g., price, static power consumption, area), as given in Table 6.1.
- (c) *Initial architecture* which is allocated based on the designers knowledge.

After parsing the above introduced files and establishing the necessary data structures, LOPOCOS applies the proposed algorithms for mapping, scheduling, and dynamic voltage scaling, which have been introduced in Chapters 3, 4, and 5. Once the optimisations have been finished, an output file is produced that contains the following information:

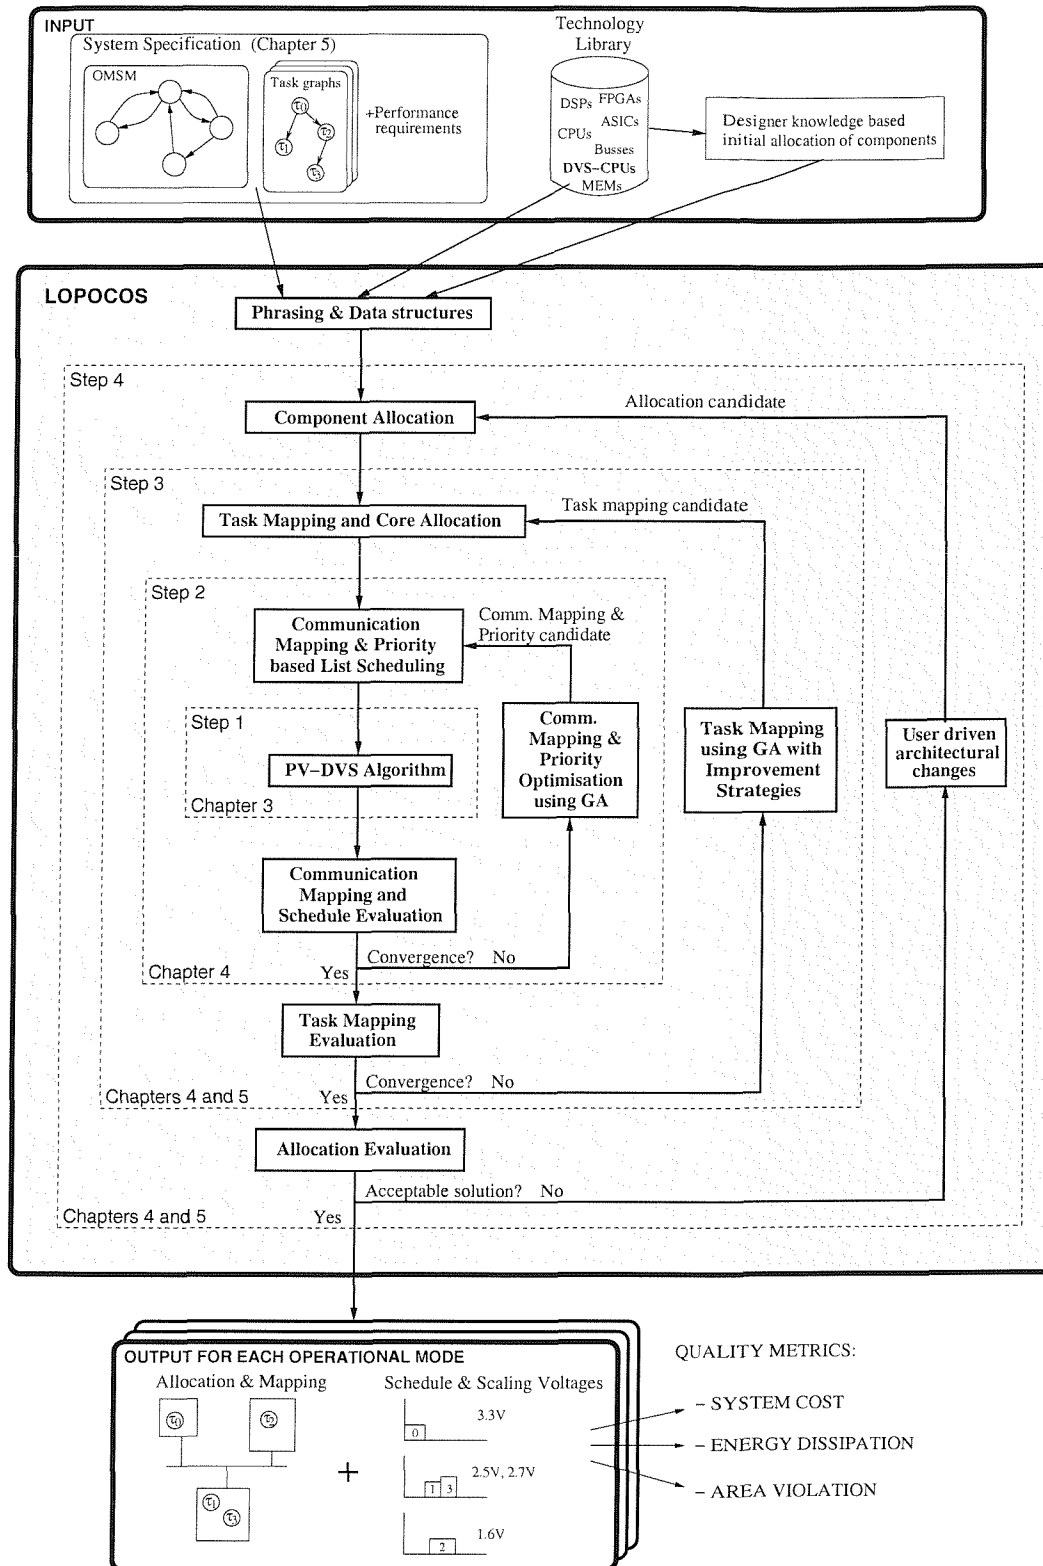


Figure 6.7: Design flow used within LOPOCOS

- (a) Mapping of tasks and communications to the active components.
- (b) Schedule for the system activities.
- (c) Scaled supply voltage for tasks executing on DVS-enabled processing elements.
- (d) Dynamic and static power consumptions.
- (e) Design quality metrics in terms of cost, area violations, and timing penalties.
- (f) General information about the synthesis run, such as number of timing infeasible and area infeasible solutions produced during the optimisation.

In the case of multi-mode systems, this output is provided for each operational mode. Furthermore, LOPOCOS returns the above information as a set of Pareto optimal solutions, i.e., several not dominated solutions with different area/energy trade-offs from which the designer can choose the most suitable design.

In order to demonstrate the tool-assisted architectural design space exploration, the remainder of this section is split into two subsections. Section 6.2.1 describes briefly the necessary input files and goes into the details of the library organisation. Section 6.2.2 outlined the usage of LOPOCOS for architectural exploration.

6.2.1 Input Descriptions

Before the system-level synthesis process can be started, it is necessary to establish the system specification as well as the technology library. The organisation of these files is explained next.

System Specification

Within LOPOCOS the system behaviour is modelled as an operational mode state machine (OMSM). That is, the overall system functionality is represented as top-level finite state machine, and each state is further described as a set of task graphs (as discussed in Chapter 5). Note, in the case of a single mode system, the OMSM consists of a single state and a single task graph. However, the OMSM for the smart phone example models the interaction of eight operational modes with 20 possible transitions (refer to Figure 5.1,

page 105). The file description of this OMSM is given in Figure 6.8, which shows the 20 mode transitions in the lines 3–22. Each of these entries corresponds to a single mode

```

1 # Possible mode transitions
2 # FromMode      ToMode      Maximal Allowed Transition Time
3 (0      --->      (1      100
4 (1      --->      (0      10
5 (1      --->      (2      100
6 (1      --->      (3      200
7 (1      --->      (4      200
8 (1      --->      (5      200
9 (2      --->      (1      50
10 (2      --->      (3      200
11 (2      --->      (6      200
12 (2      --->      (7      200
13 (3      --->      (4      50
14 (3      --->      (7      100
15 (4      --->      (1      50
16 (4      --->      (7      100
17 (5      --->      (1      50
18 (5      --->      (6      100
19 (6      --->      (2      100
20 (6      --->      (5      50
21 (7      --->      (2      100
22 (7      --->      (4      50
23
24
25 # Mode execution probabilities
26 # The sum of probabilities should equal
27 Mode0: 0.09
28 Mode1: 0.74
29 Mode2: 0.01
30 Mode3: 0.02
31 Mode4: 0.02
32 Mode5: 0.10
33 Mode6: 0.01
34 Mode7: 0.01

```

Figure 6.8: File description of the top-level finite state of the smart phone

transition and additionally states the maximal allowed time for the transition. For instance, line 7 refers to the mode transition from Radio Link Control (mode 1) to Decode Photo + RLC (mode 4), which is invoked upon a user request to restore photos, as shown in Figure 5.1. The maximal allowed time for this transition is 200ms. Furthermore, the file shown in Figure 6.8 holds the execution probability of the eight operational modes (line 27–34), towards which the design is optimised.

In order to complete the system specification, the functionality of each operational mode has to be expressed as a task graph. An example file description of a task graph is shown in Figure 6.9. This task graph represents the functionality of mode MP3 play + RLC. Since this mode fulfils two functions at the same time, MP3 decoding and radio link control, it contains two task graphs (lines 6–42 and line 47). The first graph represents the

```

1  HYPERPERIOD 0.025
2  TOLERABLE_TIMING_PENALTY 1.1
3
4  # =====MP3-DECODER=====
5
6  Task: ( 0 0 )      ttype: 46      epst: 0      dtype: NON      Deadline: 0
7  Task: ( 0 1 )      ttype: 47      epst: 0      dtype: NON      Deadline: 0
8  Task: ( 0 2 )      ttype: 47      epst: 0      dtype: NON      Deadline: 0
9  Task: ( 0 3 )      ttype: 40      epst: 0      dtype: NON      Deadline: 0
10 Task: ( 0 4 )      ttype: 40      epst: 0      dtype: NON      Deadline: 0
11 Task: ( 0 5 )      ttype: 41      epst: 0      dtype: NON      Deadline: 0
12 Task: ( 0 6 )      ttype: 41      epst: 0      dtype: NON      Deadline: 0
13 Task: ( 0 7 )      ttype: 48      epst: 0      dtype: NON      Deadline: 0
14 Task: ( 0 8 )      ttype: 49      epst: 0      dtype: NON      Deadline: 0
15 Task: ( 0 9 )      ttype: 49      epst: 0      dtype: NON      Deadline: 0
16 Task: ( 0 10 )     ttype: 50      epst: 0      dtype: NON      Deadline: 0
17 Task: ( 0 11 )     ttype: 50      epst: 0      dtype: NON      Deadline: 0
18 Task: ( 0 12 )     ttype: 43      epst: 0      dtype: NON      Deadline: 0
19 Task: ( 0 13 )     ttype: 43      epst: 0      dtype: NON      Deadline: 0
20 Task: ( 0 14 )     ttype: 51      epst: 0      dtype: NON      Deadline: 0
21 Task: ( 0 15 )     ttype: 51      epst: 0      dtype: NON      Deadline: 0
22 Task: ( 0 16 )     ttype: 52      epst: 0      dtype: HARD      Deadline: 0.025
23
24
25 Edge: ( 0 0 ) --> ( 0 1 )      etype: 6
26 Edge: ( 0 0 ) --> ( 0 2 )      etype: 6
27 Edge: ( 0 1 ) --> ( 0 3 )      etype: 6
28 Edge: ( 0 2 ) --> ( 0 4 )      etype: 6
29 Edge: ( 0 3 ) --> ( 0 5 )      etype: 7
30 Edge: ( 0 4 ) --> ( 0 6 )      etype: 7
31 Edge: ( 0 5 ) --> ( 0 7 )      etype: 7
32 Edge: ( 0 6 ) --> ( 0 7 )      etype: 7
33 Edge: ( 0 7 ) --> ( 0 8 )      etype: 7
34 Edge: ( 0 7 ) --> ( 0 9 )      etype: 7
35 Edge: ( 0 8 ) --> ( 0 10 )     etype: 7
36 Edge: ( 0 9 ) --> ( 0 11 )     etype: 7
37 Edge: ( 0 10 ) --> ( 0 12 )    etype: 7
38 Edge: ( 0 11 ) --> ( 0 13 )    etype: 7
39 Edge: ( 0 12 ) --> ( 0 14 )    etype: 7
40 Edge: ( 0 13 ) --> ( 0 15 )    etype: 7
41 Edge: ( 0 14 ) --> ( 0 16 )    etype: 7
42 Edge: ( 0 15 ) --> ( 0 16 )    etype: 7
43
44
45 # =====GSM-RLC=====
46
47 Task: ( 1 0 )      ttype: 33      epst: 0      dtype: HARD      Deadline: 0.025

```

Figure 6.9: File description of a single mode task graph

MP3 decoding consisting of 17 tasks and 18 edges between these tasks. This description corresponds to the specification given in Figure 1.3 on Page 7. Consider, for instance, lines 9 and 10 in Figure 6.9 which represent the Huffman decoder tasks τ_3 and τ_4 in Figure 1.3. The task type 40 (ttype) refers to a Huffman decoder, the earliest possible start time (epst) of the tasks is not specified, and the tasks have no imposed deadline (dtype:NON and Deadline:0). The edges between tasks are given in the lines 25–42. For example, line 29 describes the data dependency between tasks τ_3 and τ_5 , i.e., the data

which needs to be transferred from the Huffman decoder to the dequantisation unit. The amount of data is given by the edge type *etype*, e.g., *etype:7* refers to 256 bytes. A detailed description of the task graph file organisation is provided in Appendix B.

Technology Library

In addition to the system specification given as OMSM and task graphs, LOPOCOS requires the input of a technology library. This library contains modelling information about the available system components (processing elements and communication links). The overall goal of the co-synthesis is to appropriately select components from this library to implement the system's functionality according to the system specification. The selected components form the target architecture onto which the application is mapped and scheduled. Components modelled in LOPOCOS are general-purpose processors (GPPs), application-specific instruction-set processors (ASIPs), application-specific integrated circuits (ASICs), field-programmable gate arrays (FPGAs), communication links, and memory modules. A small size example of a technology library is shown Figure 6.10. It contains three different processing elements: an ARM7 DVS processor (lines 1–17), an ASIC in $0.35\mu\text{m}$ technology (lines 18–34), and a Xilinx Virtex II-pro FPGA (lines 35–51). Furthermore, a 16 bit wide communication bus (lines 52–58) and a standard 256kB memory module (lines 59–65) are given. These components are modelled through two data sets: *task independent component parameters* and *task dependent parameters*. The task independent parameters describe static values of components that are unaffected by the executed functions. With reference to Figure 6.10, the task independent parameters of the ARM7 processor are shown lines 5–8. According to the first three entries, the processor has a price of \$30, a static power dissipation of $800\mu\text{W}$, and a nominal operational frequency of 80MHz . A description of all task independent component parameters is given in Table 6.1. The task independent parameters can be directly derived from data sheets provided by the component or core manufacturers.

The second group of parameters (task dependent parameters) describe execution properties of tasks when implemented on a certain processing element. Consider, for instance, the task independent parameters given in line 12 of the technology library (Figure 6.10), where task type 1 refers to an FFT task. Executing this algorithm on the ARM7 processor requires 302300 clock cycles and dissipates an average power of 42mW .

```

1  # =====
2  # ==                ** ARM7 DVS **                ==
3  # =====
4  @GPP 0 {
5  # price      StPwr(uW)  freq      pins      BEvSleep  BEvIDLE  AddsMem (32b)
6  # 30          800        80000     36        0.4        0.004    4294967296
7  # DVS        Vmax       vt         CommBuffer  CommTime  CommPower  CommMem
8  # 1           3.8        1.2        1          0.271247  25000     13
9  #-----task properties-----
10 # type  ExeCyc      DynPwr(uW)  StMem      DynMem      Preem      Exable
11 # 0      48820      36000     7000       3000        0          1
12 # 1      302300     42000     8000       2000        0          1
13 # 2      124040     51000     400        400         0          1
14 # 3      328210     38000     1600       1400        0          1
15 # 4       99520     37000     400        400         0          1
16 # 5      895280     60000     3900       1700        0          1
17 }
18 # =====
19 # ==                ** AMS 0.35 **                ==
20 # =====
21 @ASIC 0 {
22 # price      StPwr(uW)  freq      pins      BEvSleep  BEvIDLE  Area(mm2)
23 # 60          10        2500      41        0.01       512      35.50
24 # DVS        Vmax       vt         CommBuffer  CommTime  CommPower  CommArea
25 # 0           2.12196   0.402969  1          0.124688  21        32
26 #-----task properties-----
27 # type  ExeCyc      DynPwr(uW)  Area      Exable
28 # 0      222        220        12.33     1
29 # 1      142        141        13.20     1
30 # 2       40         70         4.38     1
31 # 3      125        324         5.49     1
32 # 4      146        460         6.23     1
33 # 5       42         60         0.93     1
34 }
35 # =====
36 # ==                ** Xilinx XC2VP7 **                ==
37 # =====
38 @FPGA 0 {
39 # price      StPwr(uW)  freq      pins      BEvSleep  BEvIDLE  CLBs
40 # 80          127        2500      53        0.058      1759     1220
41 # DVS        Vmax       vt         CommBuffer  CommTime  CommPower  CommCLB  ReconT  ReconPwr
42 # 0           2.5        0.74      1          0.4339    42.79    213     1000    25.4241
43 #-----task properties-----
44 # type  ExeCyc      DynPwr(uW)  CLB      Exable
45 # 0      288        1820      252       0
46 # 1      160        2590      274       0
47 # 2      120         910       66        0
48 # 3      163        2620      75        1
49 # 4       89        4480     128       1
50 # 5      298        390       11        1
51 }
52 # =====
53 # ==                ** Bus 16bit **                ==
54 # =====
55 @LINK 0 {
56 # price      StPwr      freq      pins      BEvSleep  maxUser  PckSize  PckOver  aPower
57 # 10         775        33000     16        0.3        5        8        33      1881.648
58 }
59 # =====
60 # ==                ** Memory Module 256kb **                ==
61 # =====
62 @MEM 0 {
63 # price      StPwr      size      IOPwr
64 # 5          89.7391    256000    145.237
65 }

```

Figure 6.10: Technology library file

Abbrivation	Type	Explanation
price	GPP/FPGA/ASIC/CL/MEM	Hardware cost
StPwr	GPP/FPGA/ASIC/CL/MEM	Static power dissipation in μW
freq	GPP/FPGA/ASIC/CL	Maximal/Nominal operational frequency
pins	GPP/FPGA/ASIC/CL	Available or needed pin count for communication connection
BEvSleep	GPP/FPGA/ASIC	Break-even time from sleep mode (if implemented)
BEvIDLE	GPP/FPGA/ASIC	Break-even time from idle mode (if implemented)
AddsMem	GPP	Addressable memory size
Area	ASIC	Available area
CLBs	FPGA	Available configurable logic blocks
DVS	GPP/FPGA/ASIC	DVS enable flag (1-enabled/0-disabled)
Vmax	GPP/FPGA/ASIC	Nominal supply voltage (used for DVS)
Vt	GPP/FPGA/ASIC	Threshold voltage (used for DVS)
CommBuffer	GPP/FPGA/ASIC	Communication buffer available
CommTime	GPP/FPGA/ASIC	Bridge communication time
CommPower	GPP/FPGA/ASIC	Power dissipation during bridge communication
CommMem	GPP	Memory needed for bridge communication
CommArea	ASIC	Area needed for bridge communication
CommCLB	FPGA	CLBs needed for bridge communication
ReconT	FPGA	Clock cycles needed for a re-configuration
ReconPwr	FPGA	Power dissipation during re-configuration
maxUser	CL	Maximal users connected to the CL
PckSize	CL	Size of each package send over the CL
PckOver	CL	Overhead for each communication activity
aPower	CL	Average dynamic power dissipation
size	MEM	Memory size in bytes
IOPwr	MEM	Average power for memory accesses

Table 6.1: Task independent components parameters

Further, to store the task in the local memory 8000 bytes are required. In addition, 2000 bytes are allocated dynamically during run-time. The task's execution cannot be preempted. Refer to Table 6.2 for a description of the different task dependent parameters. Task dependent parameters are estimated through software and hardware profiling on realistic input data or through sophisticated estimation techniques [29, 54, 85, 138–140].

Abbreviation	Type	Explanation
type	GPP/ASIC/FPGA	refers to the task type (e.g., Huffman decoder)
ExeCyc	GPP/ASIC/FPGA	Number of clock cycles needed for execution
DynPwr	GPP/ASIC/FPGA	Dynamic power dissipation
StMem	GPP	Statically required memory
DynMem	GPP	Dynamically allocated memory
Preem	GPP	Preemption cycles (needed to store/restore registers)
Area	ASIC	Area needed for implementation
CLB	FPGA	CLBs needed for implementation
Exable	GPP/ASIC/FPGA	Feasible task/PE combination (can the task type be executed)

Table 6.2: Task dependent parameters

6.2.2 Architectural Design Space Exploration

A primary goal of the system designer is to find a suitable architecture that will be embedded together with the application software inside a new product. However, given the myriad of commercially available components and their possible combinations to distributed architectures, it is not hard to understand that a straight-forward architectural choice is not easy, if not impossible. The intention of LOPOCOS is to ease this design problem by equipping the designer with a tool that helps to effectively explore different architectural implementation through an automated design process. This section exemplifies the usage of LOPOCOS for tool-assisted design space exploration using the smart phone benchmark. Several different architectural choices are examined for their suitability in terms of cost, performance, area usage, and energy consumption. It is shown how the careful interpretation of the synthesis outcomes can guide the designer's architecture choices to identify solutions of high quality.

Let us consider the following design scenario. Proper product pricing as well as low energy consumption are key to a successful sale. In order to achieve the desired market success, the embedded computing system within the smart phone needs to comply with the following design constraints:

- the system price needs to be below \$120 and
- the average power dissipation should not exceed $1.6mW$ (for an average user).

Of course, in addition to these constraints the architecture needs to provide a sufficiently

Type	Make	Price	Properties
GPP0	ARM7DVS-25	\$20	high performance (32-bit, 25MHz, DVS)
GPP1	ARM7DVS-10	\$14	medium performance (32-bit, 10MHz, DVS)
GPP2	6052	\$5	low performance (8-bit, 6.6MHz, DVS)
GPP3	ST68000CP10	\$8	medium performance (16-bit, 10MHz, no DVS)
ASIC0	AMS 0.6 μ m	\$30	small area (55.2mm ² , no DVS)
ASIC1	AMS 0.6 μ m	\$60	large area (150mm ² , no DVS)
ASIC2	AMS 0.6 μ m	\$42	medium area (80mm ² , no DVS)
ASIC3	AMS 0.6 μ m	\$68	large area (150mm ² , DVS)
LINK0	PCI Bus	\$10	7 users, 33MHz, 32-bit wide
LINK1	I ² C Bus	\$5	1024 user, serial, 400kbps
LINK2	CAN Bus	\$7	serial, 1Mbps
MEM0	1MB EDO RAM	\$1	100MHz
MEM1	4MB EDO RAM	\$5	100MHz

Table 6.3: Components in a typical technology library

high performance to meet the timing constraints. Certainly, at a first glance 1.8mW might seem unrealistically low. For a standard Lithium-Ion battery with 3.7V and 600mAh, this power consumption would result in a run-time of approximately 58 days ($3.7V \cdot 0.6Ah / 1.6mW \approx 1388h$). However, considering that additional analogue circuitry (receiver, transmitter, audio amplifiers) as well as the display with background light consume a non-negligible amount of power, the overall battery life is reduced to several days [92]. Nevertheless, taking this into account, the digital part of the system (i.e. the embedded system) is liable for approximately 20–30% of the average power consumption, making it unavoidable to limit this power dissipation to 1.6mW or lower.

In order to find an architecture that fulfils the cost, power, and performance criteria, the designer can choose different components from a technology library that contains the processing elements and communication links shown in Table 6.3. In addition to the component names, the table gives some typical price values and component properties. Equipped with the system specification of the smart phone and the above given technology library, the designer sets out to examine different architectural choices for their suitability using LOPOCOS. The architectural design space exploration can be classified into two optimisation stages:

Performance optimisation The first goal during the design space exploration is the identification of such architectures that fulfil performance requirements of the application without penalties in area and cost.

Energy Minimisation Once viable architectural implementations have been established, the next stage in the optimisation hierarchy is the energy minimisation. Based on the design knowledge gathered during the performance optimisation, the designer aims to tune the architecture towards energy efficiency by exploiting DVS and by carefully increasing the available hardware area to allow the implementation of energy critical tasks in hardware.

The following considers both optimisation steps.

6.2.2.1 Performance Optimisation

Architecture 1:

The purpose of the performance optimisation is the identification of architectures that fulfil the computational requirements of the smart phone specification. In accordance to the design cost and performance constraints, the designer initially allocates components based on intuition or previous-design experience. In fact, many new system designs are upgrades from older products. Upgrading existing systems can significantly cut down the design costs. Nevertheless, the increasing complexity of new applications demands the re-evaluation of the architecture's suitability. For instance, a previously designed GSM cellular phone is implemented on an architecture that consists of a low performance 8-bit CPU (including a 1MB RAM) an ASIC with an usable area of $55.2mm^2$, both interconnected via an I²C bus. The resulting price of this configuration is \$41, according to the components cost given in Table 6.3. Considering the GSM transcoder only, the architecture provided sufficiently high performance and even resulted in some performance headroom. To evaluate if this performance headroom is large enough to additionally perform the MP3 decoder as well as the JPEG compression/decompression the designer applies LOPOCOS. Figure 6.11 shows the co-synthesis results for this architecture as a trade-off between average power dissipation and area penalty. Each diamond represents a timing feasible implementation, that is, an implementation that satisfies the imposed deadlines and repetition rates of the individual smart phone applications. Consider, for instance, the rightmost implementation candidate which has an average power dissipation of $2.5mW$. This solution shows an area penalty of 2.6, i.e., the available hardware ($55.2mm^2$) is exceeded by 2.6 times. As the figure illustrates, all solution candidates found violate the area constraints (area penalty > 1), hence, the allocated architecture does not provided




Figure 6.11: Co-synthesis results of Architecture 1

sufficient hardware area to accommodate the timing critical tasks of the smart phone example. Hence, to viable implement this solution, a hardware area of at least $143.5mm^2$ ($55.2mm^2 \cdot 2.6$) would be necessary. On the other hand, using a fast processor instead of the low performance 6052 CPU would allow to move some timing critical tasks in the design towards software implementations, and therefore reduce the required area.

Clearly, increasing the hardware area (larger ASIC) as well as allocating a fast processor would increase the system cost. In the following both possibilities are examined.

Architecture 2:

In the first possibility to improve the system design, the designer aims to fulfil the hardware requirements of at least $143.5mm^2$. This can be achieved by replacing the current ASIC 0 with an area of $55.2mm^2$ by an ASIC 3 which offers $150mm^2$ (Table 6.3). In this way, the area requirements of the smart phone specification should be satisfied. However, the large ASIC increases the price of the embedded system to \$79. The co-synthesis outcome of LOPOCOS for this configuration is illustrated in Figure 6.12 (indicated as Architecture 2). As expected, this architecture offers sufficient area and computational power to perform all applications in the smart phone device adequately. This is indicated

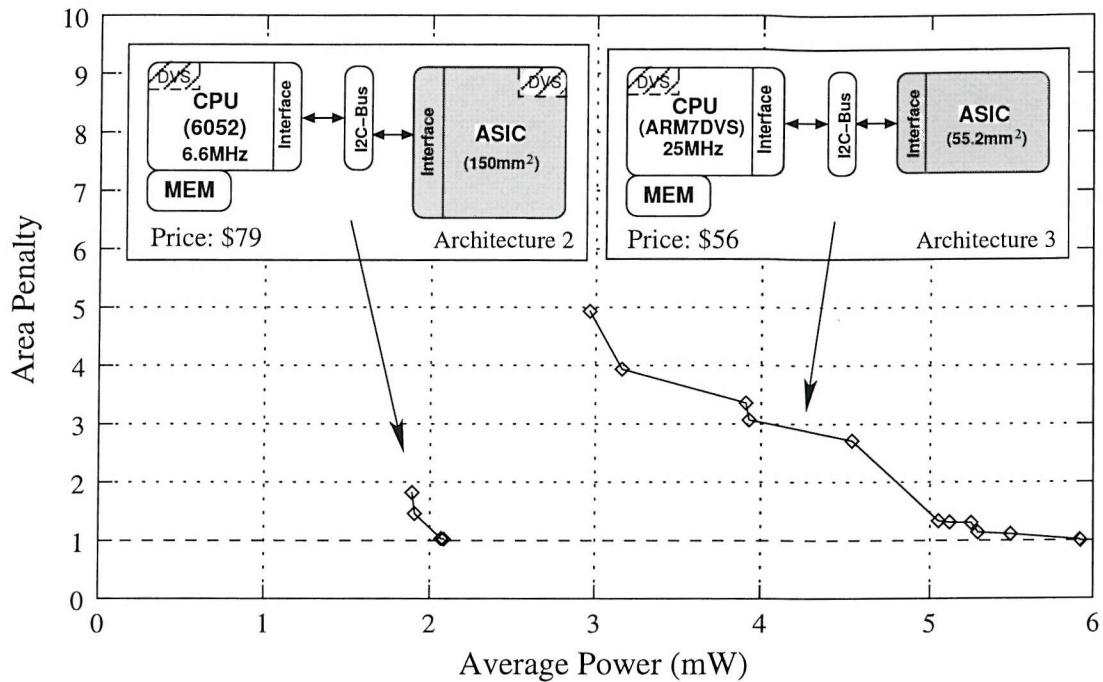


Figure 6.12: Co-synthesis results of Architectures 2 and 3

by the rightmost solution with an area penalty of 1 (no area violation) and an average power dissipation of $2.1mW$.

Architecture 3:

Let us consider the second possible modification of the initial allocation. Instead of exchanging ASIC0 with ASIC3, the 6052 8-bit CPU ($6.6MHz$) is replaced with a more powerful ARM7DVS ($25MHz$). The price for this implementation is \$56. The architecture together with the co-synthesis results are shown in Figure 6.12 (indicated as Architecture 3). As it can be observed from the figure, this combination of components allows to viably implement the applications. This is achieved due to the higher performance of the ARM7 compared to the 6052 CPU, which allows to run more of the timing critical tasks in software. Yet, the average power dissipation for the solution with no area penalty (represented rightmost diamond) is close to $6mW$, which exceeds the permitted value of $1.6mW$ by nearly 4 times. In fact, also Architecture 2, which achieved an average power dissipation of $2.1mW$ by increasing the hardware area, could not satisfy the imposed maximal power dissipation of $1.6mW$.

In summary, the design space exploration carried out so far, concentrated on area and performance aspects only. Starting from an initial yet not viable architecture (Archi-

ture 1), the design was refined in two possible directions. Firstly, by increasing the available hardware (Architecture 2), and secondly, by increase the performance of the software processor. In both case viable designs could be synthesised with LOPOCOS, however, none of which could satisfy the imposed power constraint. The aim of the following design space exploration is to further refine the system design in order to meet the power constraint.

6.2.2.2 Energy Minimisation

Having found some viable system designs (Architectures 2 and 3), the next stage in the design space exploration aims to identify solutions that comply not only timing and area constrains, but additionally limit the power consumption to imposed constraint. The main reason to perform the design space exploration for energy minimisation after the design space exploration for performance is the increased optimisation time when DVS is considered (tenth of seconds compared to 1–2 hours, as shown in the experimental results of Chapters 4 and 5). Thus, unnecessary long synthesis runs can be avoided by checking the area and timing feasibility before examining the architecture for its energy consumption.

Architectures 2 and 3:

The architectures 2 and 3 shown in the Figures 6.12 are able to adequately perform the smart phone applications. Both designs include DVS components which have not been exploited during the performance optimisation. Hence, the following examines both system designs under the consideration of DVS. Figure 6.14 shows the co-synthesis results. It can be observed that in both cases the average power dissipation was reduced, as a comparison between Figures 6.12 and 6.13 reveals. For instance, the power dissipation of Architecture 2 was reduced from $2.1mW$ to $1.9mW$. Similarly, the power dissipation of Architecture 3 was reduced from $5.9mW$ to $3.3mW$. Nevertheless, both architectures still exceed the imposed maximal power dissipation of $1.6mW$. To overcome this problem, a further design refinement is required.

Architectures 4:

In general, hardware implementations of tasks are 1–2 orders of magnitude more energy efficient than software implementations [32]. Hence, one possible way to reduce

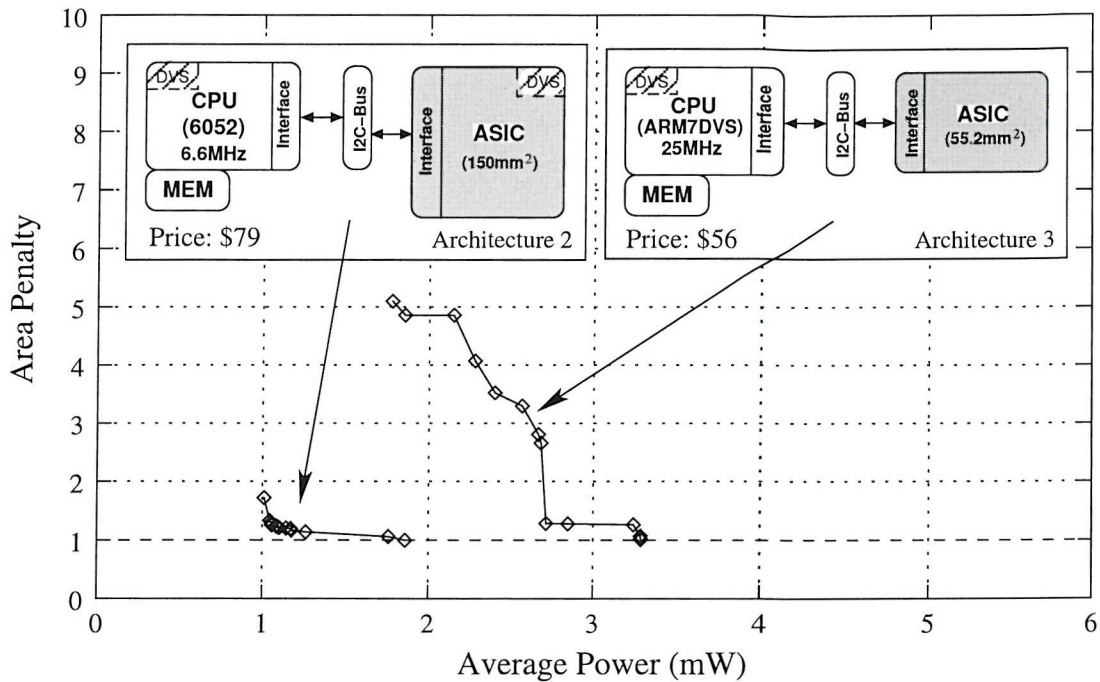


Figure 6.13: Co-synthesis results for Architectures 2 and 3, exploiting DVS

the power dissipation is map more tasks to hardware. Clearly, this necessitates to increase the available hardware area. In a similar manner, increasing the performance of software-programmable can help to reduce energy. This is achieved due to the fact that a fast processor is potentially able to satisfy the performance needs of timing critical tasks that otherwise would require a hardware implementation. Thus, more hardware area can be used to efficiently implement task with high energy dissipation. Of course, faster processor consume more power than their slower counterparts, hence the design needs to be carefully balanced. Consider, for example, the target architecture shown in Figure 6.14. This design consists of an ARM7DVS processors operating at 10MHz and two ASIC which provide areas of 150mm^2 and 55.2mm^2 . All components are connected via a single I²C bus. The total cost of this architecture is \$112, which still complies the price limit of \$120. Figure 6.14 shows the co-synthesis results for the architecture, with and without the consideration of DVS. As it can be seen from the figure, if the DVS feature of ARM7 processor is not exploited, an average power dissipation of 1.7mW can be achieved. This is still higher than the imposed power constraint of 1.6mW . However, optimising the implementation under consideration of DVS, the average power is reduced to 0.7mW . This is the first implementation that meets all imposed design constraints. The cost is below \$120, the timing constraints and area limitations are not violated, and the dissipated




Figure 6.14: Co-synthesis results for Architecture 4

average power is $0.7mW$ when using DVS.

It should be noted that this is only one possible solution. Nevertheless, the system designer might aim to find even better solutions by further exploring the architectural design space. Finding the presented solution took approximately 4.5 hours on a PentiumIII/1.2GHz PC.

6.3 Concluding Remarks

The architectural design space of distributed embedded systems is vast. Effective design tools are essential to support the system designers in exploring this solution space quickly and thoroughly. This chapter has introduced LOPOCOS, a prototype co-synthesis tool that can be used for high-level system design as a product planning tool. Considering a real-life smart phone example, it was demonstrated how LOPOCOS can be used for architectural design space exploration given a set of typical components. The aim of the design space exploration is the identification of implementation candidates that not only respect the imposed timing and area constraints of the application, but further keeps the system cost and the energy consumption within given limits. Two optimisation stages are

used to refine the design towards performance and energy minimisation. During these stages, LOPOCOS provides valuable feedback in terms of the co-synthesis trajectory. Based on this feedback, possible bottlenecks in the design are identified. This, in turn, enables the system designer to carry out necessary modifications in order to improve the system design. Using the smart phone benchmark, it was shown that potential architecture can be quickly identified.

Chapter 7

Conclusion

It is very likely that the demand for energy-efficient portable systems, as well as their complexity, particularly in consumer electronics, will continue to increase. Computational intensive applications such as multimedia with advanced audio and video coding techniques, which, until recently, were only practicable on high-performance stationary workstations, are entering the mobile arena. Designing energy-efficient embedded computing systems for such applications is a challenging and difficult task. This is because of the tight energy budget of the powering batteries as well as the massive pressure of the consumer electronics market with shortening design cycles and low product costs constraints which are essential for the success of these emerging and next-generation products. The work presented in this thesis has focused on the design of energy-efficient single-mode and multi-mode distributed embedded systems. To achieve energy-efficient designs, novel techniques and algorithms have been developed that can be used within co-synthesis frameworks to automated the design process of such embedded systems. In particular three key issues have been addressed:

- (a) *Dynamic voltage scaling* was investigated in the context of single-mode distributed heterogeneous embedded computing systems. For this purpose a number of algorithms have been proposed. A new power variation-driven voltage scaling technique, based on a novel energy-gradient strategy, which overcomes restrictions associated with traditional approaches, has been introduced. New co-synthesis techniques for scheduling and mapping have been presented, which were specifically developed for an effective utilisation of DVS by carefully increasing available slack time on DVS-enabled processing elements.

- (b) The design of *energy-efficient multi-mode embedded computing systems* has been investigated for the first time and a novel co-synthesis methodology, which enables an effective sharing of limited hardware resources by considering mode execution probabilities, has been introduced.
- (c) LOPOCOS, a prototype co-synthesis tool that incorporates the techniques and algorithms proposed in this thesis, has been developed, with the aim to extensively automate the system-level design process of energy-efficient embedded system. The automation of the design process enables an effective architectural design space exploration, which supports the system designer in finding a target architecture for high quality system implementations.

The following Section 7.1 will summarise the main contributions made by the presented work, and Section 7.2 will outline some relevant areas for future research.

7.1 Summary and Research Contributions

Dynamically voltage-scalable processors, which are capable of rapidly changing their operational state in terms of voltage and frequency setting "on-the-fly" in order to trade off energy against performance, have recently become available. The work presented in this thesis has focused on the energy minimisation of such distributed embedded systems that contain DVS processors.

Chapter 3 has introduced a new power variation-driven voltage scaling technique which takes into account the individual power dissipated by tasks during the voltage selection, in order to increase the achievable energy savings. A energy-gradient based heuristic has been developed for this purpose, which uses a local energy measurement, the energy-gradient, to scale tasks that potentially lead to high energy-efficiency. Using a mapped-and-scheduled task graph structure (MSTG) allows a propagation of execution time changes throughout the schedule in linear time; important for a fast execution of the scaling algorithm. Furthermore, the influence of a minimal extension time was investigated with regard to energy reduction possibility and performance. Experimental results have shown that up to 38% savings can be achieved for small real-life benchmarks compared to a more simple and restricted approach that uses a fixed power model.

In order to further reduce the energy dissipation of embedded systems, a system-level co-synthesis technique particular tailored for DVS was introduced in Chapter 4. By appropriately mapping and scheduling the system tasks and communications, the energy reduction capability of DVS is effectively exploited. This is achieved through a thorough mapping and scheduling optimisation based on a two-step genetic algorithm. In the first step, the schedule of system activities (tasks and communications) as well as the communication mapping are optimised not only towards timing feasibility but additionally towards energy minimisation through DVS. This is done using a genetic algorithm with a specialised genome representation that combines schedule and communication mapping into a single string. The second step is responsible for the mapping of tasks onto the processing elements of the target architecture, carefully balancing the design between area feasibility and energy minimisation. Both optimisation steps are guided by an energy estimation based on the voltage scaling technique introduced in Chapter 3. In this two-step approach, task mapping and communication mapping have been separated to avoid creation of infeasible communication mappings. As a result, the search space is restricted to structurally feasible mapping solutions only, reducing the optimisation time while maintaining the full possible optimisation potential. Extensive experiments have been conducted including a real-life benchmark of an optical flow detection algorithm. These experiments have shown that appropriate mapping and scheduling is essential in order to achieve high energy savings and simultaneously timing feasible solutions. The proposed two-step approach based on genetic algorithms accomplished this in moderate optimisation times.

Chapters 3 and 4 focused on the energy minimisation of single-mode systems. The novel co-synthesis approach introduced in Chapter 5 addressed the design of energy-efficient distributed embedded systems that need to perform multi-modes. Although many embedded systems often perform single specific functionalities, such a standalone MP3 player, modern and next-generation mobile devices demand the support for multiple operational modes, i.e., the operation across a set of several different applications. The fundamental difference between single mode and multi-mode systems is that multi-mode systems have the possibility to share computational resources among the individual modes of operation. Thereby, an effective system-level co-synthesis technique, which accounts for resource sharing, can help to significantly reduce hardware cost. The multi-mode co-synthesis approach introduced in Chapter 5 is based on a new system specification

model called operational mode state machine (OMSM), which captures mode interaction together with mode functionality. Considerable energy saving are achieved by taking into account the execution probabilities of the different modes during the co-synthesis process. These savings are possible even without the employment of DVS. The key contribution within this synthesis process was the development of a new mapping strategy, which considers the uneven execution probabilities. This mapping strategy allows the multiple task implementation on different processing elements, whenever such multiple implementations reduce the energy consumption without significantly increasing the system cost. The mapping optimisation itself is based on a genetic algorithm. Four new improvement strategies help to improve this optimisation process in terms of run-time and solution quality. This is achieved by carefully pushing the optimisation into promising search space regions. Additionally, dynamic voltage scaling was investigated in the context of multi-mode system. A virtual transformation of tasks executing on DVS-enabled hardware was shown to be useful in the presence of simultaneously scalable cores which potentially execute tasks in parallel. Overall, numerous experiments clearly indicate the advantage of taking mode execution probabilities into account during the co-synthesis.

Finally, Chapter 6 has introduced LOPOCOS, a prototype system-level co-synthesis tool which incorporates the techniques and algorithms proposed in Chapters 3, 4, and 5. The tool aids the system designer in finding system implementations that fulfil the imposed design constraints. Using a real-life smart phone example, a combination of cellular phone, MP3 player, and digital camera, it was outlined how LOPOCOS can help to effectively explore the architectural design space with the aim of finding area and timing feasible implementation that minimise system cost and energy consumption, simultaneously.

In conclusion, the broad aim of this thesis was the development of automated design techniques for energy-efficient distributed embedded systems. A detailed investigation into dynamic voltage scaling for distributed embedded systems has shown that substantial energy savings can be achieved through carefully ordering the execution of activities as well as through an appropriate mapping of the application on the target architecture. Furthermore, it was shown that the consideration of the mode execution probabilities (activation times) during the co-synthesis of multi-mode embedded systems is essential for an appropriate resource sharing that yields energy-efficient designs. Considering a realistic smart phone example it was demonstrated that up to 66% in energy savings can be

achieved compared to a naive approach that neglects execution probabilities as well as dynamic voltage scaling. In essence, when seeking to design energy-efficient distributed embedded systems, system-level co-synthesis techniques that consider energy management techniques and mode execution probabilities (as the approaches presented in this thesis) should be given serious consideration.

7.2 Future Research Directions

The present project could be elaborated in many different directions. The following introduces some interesting and relevant areas of future research.

7.2.1 Control-flow Intensive Applications

Within this work the applications that have been addressed are considered to be data-flow dominated. Such applications (e.g. voice encoding and decoding, image and video processing) can be accurately modelled using a task graph representation (see Appendix A) with the opportunity to express limited control-flow inside each task specification. However, for application with an extensive, global control-flow, models such the control process graphs (CPG) [47, 151] represent an adequate choice, since they allow a more precise and thorough capturing of this kind of system behaviour. In the case of CPGs the uncertainty about the ahead-in-time task execution represents an interesting and challenging problem, particularly when considering dynamic voltage scaling. In this context, mapping and scheduling techniques have to be developed that are specifically suitable for these system specifications.

7.2.2 Energy-aware Granularity Selection

The level of granularity has an influence on the performance optimisation potential of the application as well as on the optimisation time of the co-synthesis process [71]. Nevertheless, when designing energy-efficient systems the selected granularity can have an additional effect on the energy aspects of the implementation possibilities. Consider, for instance, a single task that can be split into two separate smaller tasks. One of which is active for most of the computations of the initial task, while the other requires only short

execution times. Clearly, when optimising the mapping it is likely that the initial (single) task is placed into hardware due to the lower execution time and power reduction (correspondingly, low energy consumption). However, in the case of the split task version, a refined mapping can place the task with high computational overhead into hardware, while the less computationally expensive task part is placed into software, avoiding unnecessary waste of hardware resources. Accordingly, the development of an energy-aware granularity selection scheme could further improve the co-synthesis process.

7.2.3 Platform Identification

As the complexity of embedded systems and underlying hardware continues to increase, the design cost for these have been risen drastically. For instance, a single mask set for a state-of-the-art chip exceed a cost of \$0.5M. With shrinking feature size this cost will further increase [6, 113]. To cope this tremendous cost the design of next-generation system-on-a-chip (SoC) embedded systems will be most likely based on platforms, i.e., pre-designed hardware architectures that can be used over a set of related applications [77]. Even if this hardware might be over-designed for the particular application at hand. Nevertheless, a challenging problem is the identification of platforms that allow a most cost-effective implementation of several different applications. Not only cost-effectiveness is essential here. Also energy consumption is an aspect that needs careful consideration when seeking to design a widely usable platform. System-level co-synthesis can help during the difficult choice of an appropriate platform. For instance, a design team might be confronted with the design of three different, yet related application (e.g., a smart phone, a PDA, and a portable MP3/DVD player). The development time and cost of these products could be significantly reduced if one would find a single hardware architecture suitable for all three applications. This platform might be over-designed to a certain extend depending on the application that needs to be implemented, but the expensive and time-consuming hardware design is reduced significantly to a single design process. Here investigations into concurrent co-design methodology for several different applications are needed. Ultimately, the development of novel "platform-level" co-synthesis techniques that consider emerging technologies such as networks-on-a-chip can aid the designers in finding a good hardware platform.

Appendix A

Supplementary Experiments

An initial assumption made in this thesis is that the execution time of tasks in transformational data-flow dominated applications is *relatively* independent of the input data stream. The following supplementary experiments are used to justify this claim. For the purpose of the presented experiments, the MP3 decoder [68] was tested using three different real-world music pieces, all consisting of approximately 3 minutes and 50 seconds of audio information, sampled at a rate of 112kbps. For all music streams the MP3 player was software profiled to extract the relevant loop counts and the function execution times. This was achieved by compiling the application with profile information [1]. After running the MP3 decoder for the three different music streams, GPROF [3] was used to extract the produced profile information. The following three listings (Tables A.1–A.3) report on the made findings. In précis, the tables show for each function within the MP3 decoder, the loop count as well as the approximate execution time information. Comparing these values, it can be observed that there is only a small deviation in the execution times of function between the different decoding runs. Furthermore, when comparing the overall decoding times (30.68s, 30.21s, and 30.23s), a maximal variation of 1.4% is given. These values can be considered as nearly static. Hence, these experiments have demonstrated that the execution times are not significantly influenced by different real-world input streams. Similar observations hold for the GSM transcoder as well as for the JPEG codec.

Function	Count	Total
ConvertToIeeeExtended	1	24.200
DetermineByteOrder	1	24.200
GetArguments	1	24.200
III_antialias	32248	22.060
III_dequantize_sample	32248	5.760
III_get_scale_factors	32248	24.160
III_get_side_info	8062	24.130
III_huffman_decode	32248	20.320
III_hybrid	1031936	23.380
III_reorder	32248	23.770
III_stereo	16124	21.560
I_CRC_calc	3	24.200
I_decode_bitalloc	3	24.200
I_decode_scale	3	24.200
Letext	0	24.200
SubBandSynthesis	580464	16.460
SwapBytesInWords	2267	23.650
Write16BitsHighLow	22	24.200
Write32BitsHighLow	10	24.200
WriteBytes	1	24.200
WriteIeeeExtendedHighLow	1	24.200
aiff_seek_to_sound_data	1	24.200
aiff_write_headers	1	24.200
alloc_buffer	1	24.200
buffer_CRC	3	24.200
close_bit_stream_r	1	24.200
create_syn_filter	1	24.200
decode_info	8065	24.200
desalloc_buffer	1	24.200
end_bs	18148	24.200
get1bit	246666	24.100
getbits	5134020	22.960
hdr_to_frps	8065	24.190
hget1bit	36835673	19.570
hgetbits	632870	24.150
hputbuf	4764148	23.950
hsstell	628651	24.070
huffman_decoder	6851927	18.560
initialize_huffman	32248	24.170
inv_mdct	1031936	11.200
main	0	21.050
main_data_slots	16124	24.200
mem_alloc	7	24.200
open_bit_stream_r	1	24.200
out_fifo	16164	22.540

Table A.1: Profiling results of music piece I

Function	Count	Total
ConvertToIeeeExtended	1	23.660
DetermineByteOrder	1	23.660
GetArguments	1	23.660
III_antialias	32188	22.100
III_dequantize_sample	32188	5.560
III_get_scale_factors	32188	23.590
III_get_side_info	8047	23.660
III_huffman_decode	32188	19.720
III_hybrid	1030016	21.590
III_reorder	32188	23.330
III_stereo	16094	22.980
I_CRC_calc	3	23.660
I_decode_bitalloc	3	23.660
I_decode_scale	3	23.660
Letext	0	23.660
SubBandSynthesis	579384	5.610
SwapBytesInWords	2263	23.190
Write16BitsHighLow	22	23.660
Write32BitsHighLow	10	23.660
WriteBytes	1	23.660
WriteIeeeExtendedHighLow	1	23.660
aiff_seek_to_sound_data	1	23.660
aiff_write_headers	1	23.660
alloc_buffer	1	23.660
buffer_CRC	3	23.660
close_bit_stream_r	1	23.660
create_syn_filter	1	23.660
decode_info	8050	23.650
desalloc_buffer	1	23.660
end_bs	18097	23.660
get1bit	246555	23.640
getbits	5124827	21.050
hdr_to_frps	8050	23.660
hget1bit	36226285	18.990
hgetbits	704216	23.610
hputbuf	4755284	23.400
hsstell	574764	23.560
huffman_decoder	6673700	17.900
initialize_huffman	32188	23.660
inv_mdct	1030016	15.980
main	0	20.430
main_data_slots	16094	23.660
mem_alloc	7	23.660
open_bit_stream_r	1	23.660
out_fifo	16134	22.550

Table A.2: Profiling results of music piece II

Function	Count	Total
ConvertToIeeeExtended	1	24.930
DetermineByteOrder	1	24.930
GetArguments	1	24.930
III_antialias	32212	24.380
III_dequantize_sample	32212	6.270
III_get_scale_factors	32212	24.920
III_get_side_info	8053	24.930
III_huffman_decode	32212	21.210
III_hybrid	1030784	23.070
III_reorder	32212	24.690
III_stereo	16106	23.990
I_CRC_calc	3	24.930
I_decode_bitalloc	3	24.930
I_decode_scale	3	24.930
SubBandSynthesis	579816	11.880
SwapBytesInWords	2264	24.630
Write16BitsHighLow	22	24.930
Write32BitsHighLow	10	24.930
WriteBytes	1	24.930
WriteIeeeExtendedHighLow	1	24.930
aiff_seek_to_sound_data	1	24.930
aiff_write_headers	1	24.930
alloc_buffer	1	24.930
buffer_CRC	3	24.930
close_bit_stream_r	1	24.930
create_syn_filter	1	24.930
decode_info	8056	24.930
desalloc_buffer	1	24.930
end_bs	18106	24.930
get1bit	246189	24.890
getbits	5128098	23.540
hdr_to_frps	8056	24.930
hget1bit	36382548	20.470
hgetbits	692573	24.790
hputbuf	4758830	24.860
hsstell	657965	24.930
huffman_decoder	6756377	19.030
initialize_huffman	32212	24.910
inv_mdct	1030784	16.820
main	0	21.950
main_data_slots	16106	24.930
mem_alloc	7	24.930
open_bit_stream_r	1	24.930
out_fifo	16146	22.520
read_syn_window	1	24.930

Table A.3: Profiling results of music piece III

Appendix B

Smart Phone Benchmark

This appendix details the system specification of the smart phone example, which contains a GSM cellular phone voice encoder and decoder, an MP3 decoder, as well as an JPEG decoder and encoder. The original C code of these algorithms can be found in [4, 5, 68]. The presented task graph representations of these algorithms were derived from these sources with slight modifications to express application parallelism. Furthermore, data dependencies between tasks was restricted to the essential data only, i.e., trivial small data flow between task were inherited into the global data flow to avoid unnecessary high communication overhead.

B.1 GSM Voice Compression

This function is responsible for voice encoding, i.e., for audio compression of human voice based on Regular Pulse Excited (RPE) coding with Long Term Prediction (LTP). A excellent description of the GSM voice encoding is provided in [69]. The following listing describes the task graph of the GSM decoder as used as input to LOPOCOS. This description corresponds to the task graph shown in Figure 6.2, page 139.

```
1  # THIS IS ALREADY THE HYPER TASK GRAPH !
2  HYPERPERIOD 0.02
3  TOLERABLE_TIMING_PENALTY 1.0
4
5  #=====GSM-ENCODER=====
6
7  Task: ( 1 0 )      ttype: 13      epst: 0      dtype: NON      Deadline: 0
8  Task: ( 1 1 )      ttype: 14      epst: 0      dtype: NON      Deadline: 0
9  Task: ( 1 2 )      ttype: 15      epst: 0      dtype: NON      Deadline: 0
10 Task: ( 1 3 )      ttype: 16      epst: 0      dtype: NON      Deadline: 0
```

```

11 Task: ( 1 4 )      ttype: 17      epst: 0      dtype: NON      Deadline: 0
12 Task: ( 1 5 )      ttype: 18      epst: 0      dtype: NON      Deadline: 0
13 Task: ( 1 6 )      ttype: 7       epst: 0      dtype: NON      Deadline: 0
14 Task: ( 1 7 )      ttype: 20      epst: 0      dtype: NON      Deadline: 0
15 Task: ( 1 8 )      ttype: 9       epst: 0      dtype: NON      Deadline: 0
16 Task: ( 1 9 )      ttype: 22      epst: 0      dtype: NON      Deadline: 0
17 Task: ( 1 10 )     ttype: 55      epst: 0      dtype: NON      Deadline: 0
18 Task: ( 1 11 )     ttype: 9       epst: 0      dtype: NON      Deadline: 0
19 Task: ( 1 12 )     ttype: 59      epst: 0      dtype: NON      Deadline: 0
20 Task: ( 1 13 )     ttype: 20      epst: 0      dtype: NON      Deadline: 0
21 Task: ( 1 14 )     ttype: 9       epst: 0      dtype: NON      Deadline: 0
22 Task: ( 1 15 )     ttype: 22      epst: 0      dtype: NON      Deadline: 0
23 Task: ( 1 16 )     ttype: 56      epst: 0      dtype: NON      Deadline: 0
24 Task: ( 1 17 )     ttype: 9       epst: 0      dtype: NON      Deadline: 0
25 Task: ( 1 18 )     ttype: 52      epst: 0      dtype: NON      Deadline: 0
26 Task: ( 1 19 )     ttype: 23      epst: 0      dtype: NON      Deadline: 0
27 Task: ( 1 20 )     ttype: 24      epst: 0      dtype: NON      Deadline: 0
28 Task: ( 1 21 )     ttype: 25      epst: 0      dtype: NON      Deadline: 0
29 Task: ( 1 22 )     ttype: 26      epst: 0      dtype: NON      Deadline: 0
30 Task: ( 1 23 )     ttype: 27      epst: 0      dtype: NON      Deadline: 0
31 Task: ( 1 24 )     ttype: 28      epst: 0      dtype: NON      Deadline: 0
32 Task: ( 1 25 )     ttype: 3       epst: 0      dtype: NON      Deadline: 0
33 Task: ( 1 26 )     ttype: 4       epst: 0      dtype: NON      Deadline: 0
34 Task: ( 1 27 )     ttype: 31      epst: 0      dtype: NON      Deadline: 0
35 Task: ( 1 28 )     ttype: 24      epst: 0      dtype: NON      Deadline: 0
36 Task: ( 1 29 )     ttype: 25      epst: 0      dtype: NON      Deadline: 0
37 Task: ( 1 30 )     ttype: 26      epst: 0      dtype: NON      Deadline: 0
38 Task: ( 1 31 )     ttype: 27      epst: 0      dtype: NON      Deadline: 0
39 Task: ( 1 32 )     ttype: 28      epst: 0      dtype: NON      Deadline: 0
40 Task: ( 1 33 )     ttype: 3       epst: 0      dtype: NON      Deadline: 0
41 Task: ( 1 34 )     ttype: 4       epst: 0      dtype: NON      Deadline: 0
42 Task: ( 1 35 )     ttype: 31      epst: 0      dtype: NON      Deadline: 0
43 Task: ( 1 36 )     ttype: 24      epst: 0      dtype: NON      Deadline: 0
44 Task: ( 1 37 )     ttype: 25      epst: 0      dtype: NON      Deadline: 0
45 Task: ( 1 38 )     ttype: 26      epst: 0      dtype: NON      Deadline: 0
46 Task: ( 1 39 )     ttype: 27      epst: 0      dtype: NON      Deadline: 0
47 Task: ( 1 40 )     ttype: 28      epst: 0      dtype: NON      Deadline: 0
48 Task: ( 1 41 )     ttype: 3       epst: 0      dtype: NON      Deadline: 0
49 Task: ( 1 42 )     ttype: 4       epst: 0      dtype: NON      Deadline: 0
50 Task: ( 1 43 )     ttype: 31      epst: 0      dtype: NON      Deadline: 0
51 Task: ( 1 44 )     ttype: 24      epst: 0      dtype: NON      Deadline: 0
52 Task: ( 1 45 )     ttype: 25      epst: 0      dtype: NON      Deadline: 0
53 Task: ( 1 46 )     ttype: 26      epst: 0      dtype: NON      Deadline: 0
54 Task: ( 1 47 )     ttype: 27      epst: 0      dtype: NON      Deadline: 0
55 Task: ( 1 48 )     ttype: 28      epst: 0      dtype: NON      Deadline: 0
56 Task: ( 1 49 )     ttype: 3       epst: 0      dtype: NON      Deadline: 0
57 Task: ( 1 50 )     ttype: 4       epst: 0      dtype: NON      Deadline: 0
58 Task: ( 1 51 )     ttype: 31      epst: 0      dtype: NON      Deadline: 0
59 Task: ( 1 52 )     ttype: 32      epst: 0      dtype: HARD      Deadline: 0.02
60
61
62 Edge: ( 1 0 ) --> ( 1 1 )      etype: 1
63 Edge: ( 1 0 ) --> ( 1 7 )      etype: 0
64 Edge: ( 1 0 ) --> ( 1 10 )     etype: 0
65 Edge: ( 1 0 ) --> ( 1 13 )     etype: 0
66 Edge: ( 1 0 ) --> ( 1 16 )     etype: 0
67 Edge: ( 1 0 ) --> ( 1 9 )      etype: 8
68 Edge: ( 1 0 ) --> ( 1 12 )     etype: 9
69 Edge: ( 1 0 ) --> ( 1 15 )     etype: 8
70 Edge: ( 1 0 ) --> ( 1 18 )     etype: 10
71 Edge: ( 1 1 ) --> ( 1 2 )      etype: 2
72 Edge: ( 1 1 ) --> ( 1 9 )      etype: 0
73 Edge: ( 1 1 ) --> ( 1 20 )     etype: 10
74 Edge: ( 1 1 ) --> ( 1 52 )     etype: 0
75 Edge: ( 1 2 ) --> ( 1 3 )      etype: 3
76 Edge: ( 1 3 ) --> ( 1 4 )      etype: 0
77 Edge: ( 1 4 ) --> ( 1 5 )      etype: 0
78 Edge: ( 1 5 ) --> ( 1 6 )      etype: 0
79 Edge: ( 1 6 ) --> ( 1 7 )      etype: 0
80 Edge: ( 1 6 ) --> ( 1 10 )     etype: 0

```

```

81 Edge: ( 1 6 ) --> ( 1 13 ) etype: 0
82 Edge: ( 1 6 ) --> ( 1 16 ) etype: 0
83 Edge: ( 1 7 ) --> ( 1 8 ) etype: 0
84 Edge: ( 1 8 ) --> ( 1 9 ) etype: 0
85 Edge: ( 1 9 ) --> ( 1 12 ) etype: 0
86 Edge: ( 1 9 ) --> ( 1 19 ) etype: 8
87 Edge: ( 1 10 ) --> ( 1 11 ) etype: 0
88 Edge: ( 1 11 ) --> ( 1 12 ) etype: 0
89 Edge: ( 1 12 ) --> ( 1 15 ) etype: 0
90 Edge: ( 1 12 ) --> ( 1 19 ) etype: 9
91 Edge: ( 1 13 ) --> ( 1 14 ) etype: 0
92 Edge: ( 1 14 ) --> ( 1 15 ) etype: 0
93 Edge: ( 1 15 ) --> ( 1 18 ) etype: 0
94 Edge: ( 1 15 ) --> ( 1 19 ) etype: 8
95 Edge: ( 1 16 ) --> ( 1 17 ) etype: 0
96 Edge: ( 1 17 ) --> ( 1 18 ) etype: 0
97 Edge: ( 1 18 ) --> ( 1 19 ) etype: 10
98 Edge: ( 1 18 ) --> ( 1 52 ) etype: 0
99 Edge: ( 1 19 ) --> ( 1 20 ) etype: 11
100 Edge: ( 1 19 ) --> ( 1 28 ) etype: 11
101 Edge: ( 1 19 ) --> ( 1 36 ) etype: 11
102 Edge: ( 1 19 ) --> ( 1 44 ) etype: 11
103 Edge: ( 1 20 ) --> ( 1 21 ) etype: 12
104 Edge: ( 1 21 ) --> ( 1 22 ) etype: 13
105 Edge: ( 1 21 ) --> ( 1 27 ) etype: 11
106 Edge: ( 1 22 ) --> ( 1 23 ) etype: 11
107 Edge: ( 1 23 ) --> ( 1 24 ) etype: 8
108 Edge: ( 1 23 ) --> ( 1 26 ) etype: 16
109 Edge: ( 1 24 ) --> ( 1 25 ) etype: 17
110 Edge: ( 1 25 ) --> ( 1 26 ) etype: 9
111 Edge: ( 1 26 ) --> ( 1 27 ) etype: 14
112 Edge: ( 1 27 ) --> ( 1 28 ) etype: 10
113 Edge: ( 1 28 ) --> ( 1 29 ) etype: 12
114 Edge: ( 1 29 ) --> ( 1 30 ) etype: 13
115 Edge: ( 1 29 ) --> ( 1 35 ) etype: 11
116 Edge: ( 1 30 ) --> ( 1 31 ) etype: 11
117 Edge: ( 1 31 ) --> ( 1 32 ) etype: 8
118 Edge: ( 1 31 ) --> ( 1 34 ) etype: 16
119 Edge: ( 1 32 ) --> ( 1 33 ) etype: 17
120 Edge: ( 1 33 ) --> ( 1 34 ) etype: 9
121 Edge: ( 1 34 ) --> ( 1 35 ) etype: 14
122 Edge: ( 1 35 ) --> ( 1 36 ) etype: 10
123 Edge: ( 1 36 ) --> ( 1 37 ) etype: 12
124 Edge: ( 1 37 ) --> ( 1 38 ) etype: 13
125 Edge: ( 1 37 ) --> ( 1 43 ) etype: 11
126 Edge: ( 1 38 ) --> ( 1 39 ) etype: 11
127 Edge: ( 1 39 ) --> ( 1 40 ) etype: 8
128 Edge: ( 1 39 ) --> ( 1 42 ) etype: 16
129 Edge: ( 1 40 ) --> ( 1 41 ) etype: 17
130 Edge: ( 1 41 ) --> ( 1 42 ) etype: 9
131 Edge: ( 1 42 ) --> ( 1 43 ) etype: 14
132 Edge: ( 1 43 ) --> ( 1 44 ) etype: 10
133 Edge: ( 1 44 ) --> ( 1 45 ) etype: 12
134 Edge: ( 1 45 ) --> ( 1 46 ) etype: 13
135 Edge: ( 1 45 ) --> ( 1 51 ) etype: 11
136 Edge: ( 1 46 ) --> ( 1 47 ) etype: 11
137 Edge: ( 1 47 ) --> ( 1 48 ) etype: 8
138 Edge: ( 1 47 ) --> ( 1 50 ) etype: 16
139 Edge: ( 1 48 ) --> ( 1 49 ) etype: 17
140 Edge: ( 1 49 ) --> ( 1 50 ) etype: 9
141 Edge: ( 1 50 ) --> ( 1 51 ) etype: 14
142 Edge: ( 1 51 ) --> ( 1 52 ) etype: 10

```

B.2 GSM Voice Decompression

This function reverses the encoding process and hence restores the human voice from the 13kbps received bit stream. It corresponds to the task graph given in Figure 6.3, page 140.

```

1  # THIS IS ALREADY THE HYPER TASK GRAPH !
2  HYPERPERIOD 0.02
3  TOLERABLE_TIMING_PENALTY 1.0
4
5  # =====GSM DECODER=====
6
7  Task: ( 0 0 )      ttype: 0      epst: 0      dtype: NON      Deadline: 0
8  Task: ( 0 1 )      ttype: 1      epst: 0      dtype: NON      Deadline: 0
9  Task: ( 0 2 )      ttype: 1      epst: 0      dtype: NON      Deadline: 0
10 Task: ( 0 3 )      ttype: 1      epst: 0      dtype: NON      Deadline: 0
11 Task: ( 0 4 )      ttype: 1      epst: 0      dtype: NON      Deadline: 0
12 Task: ( 0 5 )      ttype: 2      epst: 0      dtype: NON      Deadline: 0
13 Task: ( 0 6 )      ttype: 3      epst: 0      dtype: NON      Deadline: 0
14 Task: ( 0 7 )      ttype: 4      epst: 0      dtype: NON      Deadline: 0
15 Task: ( 0 8 )      ttype: 5      epst: 0      dtype: NON      Deadline: 0
16 Task: ( 0 9 )      ttype: 3      epst: 0      dtype: NON      Deadline: 0
17 Task: ( 0 10 )     ttype: 4      epst: 0      dtype: NON      Deadline: 0
18 Task: ( 0 11 )     ttype: 5      epst: 0      dtype: NON      Deadline: 0
19 Task: ( 0 12 )     ttype: 3      epst: 0      dtype: NON      Deadline: 0
20 Task: ( 0 13 )     ttype: 4      epst: 0      dtype: NON      Deadline: 0
21 Task: ( 0 14 )     ttype: 5      epst: 0      dtype: NON      Deadline: 0
22 Task: ( 0 15 )     ttype: 3      epst: 0      dtype: NON      Deadline: 0
23 Task: ( 0 16 )     ttype: 4      epst: 0      dtype: NON      Deadline: 0
24 Task: ( 0 17 )     ttype: 5      epst: 0      dtype: NON      Deadline: 0
25 Task: ( 0 18 )     ttype: 6      epst: 0      dtype: NON      Deadline: 0
26 Task: ( 0 19 )     ttype: 7      epst: 0      dtype: NON      Deadline: 0
27 Task: ( 0 20 )     ttype: 12     epst: 0      dtype: HARD      Deadline: 0.02
28 Task: ( 0 21 )     ttype: 8      epst: 0      dtype: NON      Deadline: 0
29 Task: ( 0 22 )     ttype: 9      epst: 0      dtype: NON      Deadline: 0
30 Task: ( 0 23 )     ttype: 10     epst: 0      dtype: NON      Deadline: 0
31 Task: ( 0 24 )     ttype: 54     epst: 0      dtype: NON      Deadline: 0
32 Task: ( 0 25 )     ttype: 9      epst: 0      dtype: NON      Deadline: 0
33 Task: ( 0 26 )     ttype: 57     epst: 0      dtype: NON      Deadline: 0
34 Task: ( 0 27 )     ttype: 8      epst: 0      dtype: NON      Deadline: 0
35 Task: ( 0 28 )     ttype: 9      epst: 0      dtype: NON      Deadline: 0
36 Task: ( 0 29 )     ttype: 10     epst: 0      dtype: NON      Deadline: 0
37 Task: ( 0 30 )     ttype: 53     epst: 0      dtype: NON      Deadline: 0
38 Task: ( 0 31 )     ttype: 9      epst: 0      dtype: NON      Deadline: 0
39 Task: ( 0 32 )     ttype: 58     epst: 0      dtype: NON      Deadline: 0
40 Task: ( 0 33 )     ttype: 11     epst: 0      dtype: HARD      Deadline: 0.02
41
42 Edge: ( 0 0 ) --> ( 0 8 )      etype: 10
43 Edge: ( 0 0 ) --> ( 0 21 )     etype: 4
44 Edge: ( 0 0 ) --> ( 0 24 )     etype: 4
45 Edge: ( 0 0 ) --> ( 0 27 )     etype: 4
46 Edge: ( 0 0 ) --> ( 0 30 )     etype: 4
47 Edge: ( 0 0 ) --> ( 0 23 )     etype: 4
48 Edge: ( 0 1 ) --> ( 0 6 )      etype: 18
49 Edge: ( 0 1 ) --> ( 0 7 )      etype: 16
50 Edge: ( 0 1 ) --> ( 0 8 )      etype: 19
51 Edge: ( 0 2 ) --> ( 0 9 )      etype: 18
52 Edge: ( 0 2 ) --> ( 0 10 )     etype: 16
53 Edge: ( 0 2 ) --> ( 0 11 )     etype: 19
54 Edge: ( 0 3 ) --> ( 0 12 )     etype: 18
55 Edge: ( 0 3 ) --> ( 0 13 )     etype: 16
56 Edge: ( 0 3 ) --> ( 0 14 )     etype: 19
57 Edge: ( 0 4 ) --> ( 0 15 )     etype: 18
58 Edge: ( 0 4 ) --> ( 0 16 )     etype: 16
59 Edge: ( 0 4 ) --> ( 0 17 )     etype: 19
60 Edge: ( 0 5 ) --> ( 0 19 )     etype: 0
61 Edge: ( 0 6 ) --> ( 0 7 )      etype: 20

```

```

62 Edge: ( 0 7 ) --> ( 0 8 )      etype: 11
63 Edge: ( 0 8 ) --> ( 0 11 )     etype: 10
64 Edge: ( 0 9 ) --> ( 0 10 )     etype: 20
65 Edge: ( 0 10 ) --> ( 0 11 )    etype: 11
66 Edge: ( 0 11 ) --> ( 0 14 )    etype: 10
67 Edge: ( 0 12 ) --> ( 0 13 )    etype: 20
68 Edge: ( 0 13 ) --> ( 0 14 )    etype: 11
69 Edge: ( 0 14 ) --> ( 0 17 )    etype: 10
70 Edge: ( 0 15 ) --> ( 0 16 )    etype: 20
71 Edge: ( 0 16 ) --> ( 0 17 )    etype: 11
72 Edge: ( 0 17 ) --> ( 0 18 )    etype: 2
73 Edge: ( 0 17 ) --> ( 0 20 )    etype: 10
74 Edge: ( 0 18 ) --> ( 0 23 )    etype: 8
75 Edge: ( 0 18 ) --> ( 0 26 )    etype: 9
76 Edge: ( 0 18 ) --> ( 0 29 )    etype: 8
77 Edge: ( 0 18 ) --> ( 0 32 )    etype: 10
78 Edge: ( 0 19 ) --> ( 0 21 )    etype: 0
79 Edge: ( 0 19 ) --> ( 0 24 )    etype: 0
80 Edge: ( 0 19 ) --> ( 0 27 )    etype: 0
81 Edge: ( 0 19 ) --> ( 0 30 )    etype: 0
82 Edge: ( 0 21 ) --> ( 0 22 )    etype: 0
83 Edge: ( 0 22 ) --> ( 0 23 )    etype: 0
84 Edge: ( 0 23 ) --> ( 0 26 )    etype: 0
85 Edge: ( 0 23 ) --> ( 0 33 )    etype: 8
86 Edge: ( 0 24 ) --> ( 0 25 )    etype: 0
87 Edge: ( 0 25 ) --> ( 0 26 )    etype: 0
88 Edge: ( 0 26 ) --> ( 0 29 )    etype: 0
89 Edge: ( 0 26 ) --> ( 0 33 )    etype: 9
90 Edge: ( 0 27 ) --> ( 0 28 )    etype: 0
91 Edge: ( 0 28 ) --> ( 0 29 )    etype: 0
92 Edge: ( 0 29 ) --> ( 0 32 )    etype: 0
93 Edge: ( 0 29 ) --> ( 0 33 )    etype: 8
94 Edge: ( 0 30 ) --> ( 0 31 )    etype: 0
95 Edge: ( 0 31 ) --> ( 0 32 )    etype: 0
96 Edge: ( 0 32 ) --> ( 0 20 )    etype: 0
97 Edge: ( 0 32 ) --> ( 0 33 )    etype: 10

```

B.3 MP3 Sound Decoder

The following specification file provides the MP3 decoder that has been introduced in Section 1.2.1, shown in Figure 1.3. This decoder reads encoded audio information from an data stream, and reproduces through several lossy transformations the "original" audio signal. To ensure a high quality audio signal, a repetition rate of 40Hz should be constantly maintained.

```

1 # THIS IS ALREADY THE HYPER TASK GRAPH !
2 HYPERPERIOD 0.025
3 TOLERABLE_TIMING_PENALTY 1.1
4
5 # =====MP3-DECODER=====
6
7 Task: ( 0 0 )      ttype: 46      epst: 0      dtype: NON      Deadline: 0
8 Task: ( 0 1 )      ttype: 47      epst: 0      dtype: NON      Deadline: 0
9 Task: ( 0 2 )      ttype: 47      epst: 0      dtype: NON      Deadline: 0
10 Task: ( 0 3 )      ttype: 40      epst: 0      dtype: NON      Deadline: 0
11 Task: ( 0 4 )      ttype: 40      epst: 0      dtype: NON      Deadline: 0
12 Task: ( 0 5 )      ttype: 41      epst: 0      dtype: NON      Deadline: 0
13 Task: ( 0 6 )      ttype: 41      epst: 0      dtype: NON      Deadline: 0
14 Task: ( 0 7 )      ttype: 48      epst: 0      dtype: NON      Deadline: 0

```



```

15 Task: ( 0 8 )      ttype: 49      epst: 0      dtype: NON      Deadline: 0
16 Task: ( 0 9 )      ttype: 49      epst: 0      dtype: NON      Deadline: 0
17 Task: ( 0 10 )     ttype: 50      epst: 0      dtype: NON      Deadline: 0
18 Task: ( 0 11 )     ttype: 50      epst: 0      dtype: NON      Deadline: 0
19 Task: ( 0 12 )     ttype: 43      epst: 0      dtype: NON      Deadline: 0
20 Task: ( 0 13 )     ttype: 43      epst: 0      dtype: NON      Deadline: 0
21 Task: ( 0 14 )     ttype: 51      epst: 0      dtype: HARD      Deadline: 0.025
22 Task: ( 0 15 )     ttype: 51      epst: 0      dtype: HARD      Deadline: 0.025
23
24
25 Edge: ( 0 0 ) --> ( 0 1 )      etype: 6
26 Edge: ( 0 0 ) --> ( 0 2 )      etype: 6
27 Edge: ( 0 1 ) --> ( 0 3 )      etype: 6
28 Edge: ( 0 2 ) --> ( 0 4 )      etype: 6
29 Edge: ( 0 3 ) --> ( 0 5 )      etype: 6
30 Edge: ( 0 4 ) --> ( 0 6 )      etype: 6
31 Edge: ( 0 5 ) --> ( 0 7 )      etype: 6
32 Edge: ( 0 6 ) --> ( 0 7 )      etype: 6
33 Edge: ( 0 7 ) --> ( 0 8 )      etype: 6
34 Edge: ( 0 7 ) --> ( 0 9 )      etype: 6
35 Edge: ( 0 8 ) --> ( 0 10 )     etype: 6
36 Edge: ( 0 9 ) --> ( 0 11 )     etype: 6
37 Edge: ( 0 10 ) --> ( 0 12 )    etype: 6
38 Edge: ( 0 11 ) --> ( 0 13 )    etype: 6
39 Edge: ( 0 12 ) --> ( 0 14 )    etype: 6
40 Edge: ( 0 13 ) --> ( 0 15 )    etype: 6

```

B.4 JPEG Image Encoder

The JPEG encoder within the smart phone device is responsible for the compression of images that have been taken with the integrated digital camera. Unlike the previous functionalities, the image encoder does not need a constant repetition rate to ensure an appropriate quality. A violation of the repetition rate of 20% can be tolerated (TOLERABLE_TIMING_PENALTY 1.2). Yet, to achieve a quick image encoding within a quarter of second it is necessary to iterate the task graph in $20\mu s$. Thereby, allowing the user to take a maximum of 4 pictures (1024x768 pixels) per second.

```

1  # THIS IS ALREADY THE HYPER TASK GRAPH !
2  HYPERPERIOD 0.000020
3  TOLERABLE_TIMING_PENALTY 1.2
4
5  # =====JPEG-ENCODER=====
6
7  Task: ( 0 0 )      ttype: 35      epst: 0      dtype: NON      Deadline: 0
8  Task: ( 0 1 )      ttype: 36      epst: 0      dtype: NON      Deadline: 0
9  Task: ( 0 2 )      ttype: 37      epst: 0      dtype: NON      Deadline: 0
10 Task: ( 0 3 )      ttype: 38      epst: 0      dtype: NON      Deadline: 0
11 Task: ( 0 4 )      ttype: 39      epst: 0      dtype: HARD      Deadline: 0.000020
12
13
14 Edge: ( 0 0 ) --> ( 0 1 )      etype: 7
15 Edge: ( 0 1 ) --> ( 0 2 )      etype: 7
16 Edge: ( 0 2 ) --> ( 0 3 )      etype: 7
17 Edge: ( 0 3 ) --> ( 0 4 )      etype: 7

```

B.5 JPEG Image Decoder

The JPEG decoder is used to restore the stored images in order to display them on the backside TFT-display. This algorithm should allow the user to flick through the already taken photos, with up to two pictures per second (2Hz). Similar to the encoder, the repetition rate of the decompression can tolerate a penalty of up to 20%.

```

1  # THIS IS ALREADY THE HYPER TASK GRAPH !
2  HYPERPERIOD 0.00004
3  TOLERABLE_TIMING_PENALTY 1.2
4
5  # =====JPEG-DECODER=====
6
7  Task: ( 0 0 )      ttype: 40      epst: 0      dtype: NON      Deadline: 0
8  Task: ( 0 1 )      ttype: 41      epst: 0      dtype: NON      Deadline: 0
9  Task: ( 0 2 )      ttype: 42      epst: 0      dtype: NON      Deadline: 0
10 Task: ( 0 3 )      ttype: 43      epst: 0      dtype: NON      Deadline: 0
11 Task: ( 0 4 )      ttype: 44      epst: 0      dtype: NON      Deadline: 0
12 Task: ( 0 5 )      ttype: 45      epst: 0      dtype: HARD      Deadline: 0.00004
13
14
15 Edge: ( 0 0 ) --> ( 0 1 )      etype: 7
16 Edge: ( 0 1 ) --> ( 0 2 )      etype: 7
17 Edge: ( 0 2 ) --> ( 0 3 )      etype: 7
18 Edge: ( 0 3 ) --> ( 0 4 )      etype: 7
19 Edge: ( 0 4 ) --> ( 0 5 )      etype: 7

```

B.6 Technology File

The processing elements described in this library are based on data that has been derived from real-world processors such as the dynamic voltage scalable CPU developed by Burd *et al.* [31, 33]. As outlined in Chapter 6, this file contains the task independent and the task dependent parameters. In detail, three processors, three ASICs, and a FPGA are described. Furthermore, the smart phone specification consists of 60 different task types, each representing the functionality of a single task. The task type is used to look up the execution properties, such as execution time and dynamic power dissipation, from the technology library. Furthermore, associating a type with each task enables an easy consideration of resource sharing, since tasks of identical types can be executed on the same components. The annotation of a task with a deadline type allows the synthesis of systems that must execute mixed application with hard real-time and soft real-time constraints. This feature can be use specify soft deadline tasks, which would degrade the quality of the system when exceeded. This is unlike hard deadline tasks which have to be met during run-time. Similarly to the tasks, the communication edges between two tasks

are annotated with a type (etype). This index specifies the data amount that needs to be transferred and it is looked up in the technology library.

```

1  # THIS IS THE TECHNOLOGY FILE
2  # last modified 10th July 2002 (Processor data according to ARM and Burd)
3
4  @GPP 0 {
5  # price      StPwr      freq      pins      BEvSleep  BEvIDLE  AddsMem  DVS
6  100          800        25000    26         0.4       0.004    256000   1
7
8  # Vmax       vt          CommBuffer CommTime  CommPower  CommMem
9  1.2          0.4        1         0.271247  2000       13
10
11 # -----
12 # type version ExeCyc  DynPwr  StMem  DynMem  Preem  Exable
13 0 0 2882 25000 8000 400 1 1
14 1 0 200 25000 8100 400 1 1
15 2 0 100 25000 7300 400 1 1
16 3 0 981 25000 420 400 1 1
17 4 0 395 25000 400 400 1 1
18 5 0 2952 25000 5300 400 1 1
19 6 0 2000 25000 7600 400 1 1
20 7 0 160 25000 2900 400 1 1
21 8 0 115 25000 1400 400 1 1
22 9 0 552 25000 400 400 1 1
23 10 0 15200 25000 10500 400 1 1
24 11 0 15596 25000 20000 400 1 1
25 12 0 2882 25000 5500 400 1 1
26 13 0 3617 25000 5400 400 1 1
27 14 0 23628 25000 400 400 1 1
28 15 0 85864 25000 8400 400 1 1
29 16 0 3460 25000 2300 400 1 1
30 17 0 472 25000 900 400 1 1
31 18 0 315 25000 600 400 1 1
32 19 0 157 25000 100 400 1 1
33 20 0 52 25000 200 400 1 1
34 21 0 197 25000 200 400 1 1
35 22 0 26154 25000 1900 400 1 1
36 23 0 3617 25000 2100 400 1 1
37 24 0 26007 25000 4600 400 1 1
38 25 0 5013 25000 8600 400 1 1
39 26 0 5014 25000 3400 400 1 1
40 27 0 1140 25000 4700 400 1 1
41 28 0 1966 25000 3700 400 1 1
42 29 0 787 25000 1400 400 1 1
43 30 0 472 25000 9400 400 1 1
44 31 0 200 25000 8200 400 1 1
45 32 0 3617 25000 800 400 1 1
46 33 0 1500 25000 900 400 1 1
47 34 0 1200 25000 1900 400 1 1
48 35 0 41580 25000 10900 400 1 1
49 36 0 138666 25000 22600 400 1 1
50 37 0 180246 25000 32500 400 1 1
51 38 0 177474 25000 3200 400 1 1
52 39 0 122034 25000 1300 400 1 1
53 40 0 36781 25000 400 400 1 1
54 41 0 14172 25000 3400 400 1 1
55 42 0 38982 25000 4200 400 1 1
56 43 0 144924 25000 20400 400 1 1
57 44 0 79250 25000 2400 400 1 1
58 45 0 135890 25000 44400 400 1 1
59 46 0 1071 25000 7700 400 1 1
60 47 0 476 25000 200 400 1 1
61 48 0 63914 25000 1300 400 1 1
62 49 0 2568 25000 600 400 1 1
63 50 0 21305 25000 3300 400 1 1
64 51 0 266687 25000 10290 400 1 1
65 52 0 78462 25000 4900 400 1 1

```

```

66 53 0 1060 25000 500 400 1 1
67 54 0 123 25000 200 400 1 1
68 55 0 55 25000 100 400 1 1
69 56 0 471 25000 200 400 1 1
70 57 0 16370 25000 200 400 1 1
71 58 0 140310 25000 1400 400 1 1
72 59 0 9154 25000 5400 400 1 1
73 }
74
75
76 @GPP 1 {
77 # price StPwr freq pins BEvSleep BEvIDLE AddsMem DVS
78 50 800 10000 26 0.5 0.005 256000 1
79
80 # Vmax vt CommBuffer CommTime CommPower CommMem
81 1.8 0.6 1 0.271247 2000 13
82
83 #-----
84 # type version ExeCyc DynPwr StMem DynMem Preem Exable
85 0 0 2882 12000 8000 400 1 1
86 1 0 200 12000 8100 400 1 1
87 2 0 100 12000 7300 400 1 1
88 3 0 981 12000 420 400 1 1
89 4 0 395 12000 400 400 1 1
90 5 0 2952 12000 5300 400 1 1
91 6 0 2000 12000 7600 400 1 1
92 7 0 160 12000 2900 400 1 1
93 8 0 115 12000 1400 400 1 1
94 9 0 552 12000 400 400 1 1
95 10 0 15200 12000 10500 400 1 1
96 11 0 15596 12000 20000 400 1 1
97 12 0 2882 12000 5500 400 1 1
98 13 0 3617 12000 5400 400 1 1
99 14 0 23628 12000 400 400 1 1
100 15 0 85864 12000 8400 400 1 1
101 16 0 3460 12000 2300 400 1 1
102 17 0 472 12000 900 400 1 1
103 18 0 315 12000 600 400 1 1
104 19 0 157 12000 100 400 1 1
105 20 0 52 12000 200 400 1 1
106 21 0 197 12000 200 400 1 1
107 22 0 26154 12000 1900 400 1 1
108 23 0 3617 12000 2100 400 1 1
109 24 0 26007 12000 4600 400 1 1
110 25 0 5013 12000 8600 400 1 1
111 26 0 5014 12000 3400 400 1 1
112 27 0 1140 12000 4700 400 1 1
113 28 0 1966 12000 3700 400 1 1
114 29 0 787 12000 1400 400 1 1
115 30 0 472 12000 9400 400 1 1
116 31 0 200 12000 8200 400 1 1
117 32 0 3617 12000 800 400 1 1
118 33 0 1500 12000 900 400 1 1
119 34 0 1200 12000 1900 400 1 1
120 35 0 41580 12000 10900 400 1 1
121 36 0 138666 12000 22600 400 1 1
122 37 0 180246 12000 32500 400 1 1
123 38 0 177474 12000 3200 400 1 1
124 39 0 122034 12000 1300 400 1 1
125 40 0 36781 12000 400 400 1 1
126 41 0 14172 12000 3400 400 1 1
127 42 0 38982 12000 4200 400 1 1
128 43 0 144924 12000 20400 400 1 1
129 44 0 79250 12000 2400 400 1 1
130 45 0 135890 12000 44400 400 1 1
131 46 0 1071 12000 7700 400 1 1
132 47 0 476 12000 200 400 1 1
133 48 0 63914 12000 1300 400 1 1
134 49 0 2568 12000 600 400 1 1
135 50 0 21305 12000 3300 400 1 1

```

```

136 51 0 266687 12000 10290 400 1 1
137 52 0 78462 12000 4900 400 1 1
138 53 0 1060 12000 500 400 1 1
139 54 0 123 12000 200 400 1 1
140 55 0 55 12000 100 400 1 1
141 56 0 471 12000 200 400 1 1
142 57 0 16370 12000 200 400 1 1
143 58 0 140310 12000 1400 400 1 1
144 59 0 9154 12000 5400 400 1 1
145 }
146
147
148 @GPP 2 {
149 # price StPwr freq pins BEvSleep BEvIDLE AddsMem DVS
150 50 800 6600 26 0.24 0.04 256000 1
151
152 # Vmax vt CommBuffer CommTime CommPower CommMem
153 1.8 0.6 1 0.271247 18000 13
154
155 #-----
156 # type version ExeCyc DynPwr StMem DynMem Preem Exable
157 0 0 2882 9000 8000 400 1 1
158 1 0 200 9000 8100 400 1 1
159 2 0 100 9000 7300 400 1 1
160 3 0 981 9000 420 400 1 1
161 4 0 395 9000 400 400 1 1
162 5 0 2952 9000 5300 400 1 1
163 6 0 2000 9000 7600 400 1 1
164 7 0 160 9000 2900 400 1 1
165 8 0 115 9000 1400 400 1 1
166 9 0 552 9000 400 400 1 1
167 10 0 15200 9000 10500 400 1 1
168 11 0 15596 9000 20000 400 1 1
169 12 0 2882 9000 5500 400 1 1
170 13 0 3617 9000 5400 400 1 1
171 14 0 23628 9000 400 400 1 1
172 15 0 85864 9000 8400 400 1 1
173 16 0 3460 9000 2300 400 1 1
174 17 0 472 9000 900 400 1 1
175 18 0 315 9000 600 400 1 1
176 19 0 157 9000 100 400 1 1
177 20 0 52 9000 200 400 1 1
178 21 0 197 9000 200 400 1 1
179 22 0 26154 9000 1900 400 1 1
180 23 0 3617 9000 2100 400 1 1
181 24 0 26007 9000 4600 400 1 1
182 25 0 5013 9000 8600 400 1 1
183 26 0 5014 9000 3400 400 1 1
184 27 0 1140 9000 4700 400 1 1
185 28 0 1966 9000 3700 400 1 1
186 29 0 787 9000 1400 400 1 1
187 30 0 472 9000 9400 400 1 1
188 31 0 200 9000 8200 400 1 1
189 32 0 3617 9000 800 400 1 1
190 33 0 1500 9000 900 400 1 1
191 34 0 1200 9000 1900 400 1 1
192 35 0 41580 9000 10900 400 1 1
193 36 0 138666 9000 22600 400 1 1
194 37 0 180246 9000 32500 400 1 1
195 38 0 177474 9000 3200 400 1 1
196 39 0 122034 9000 1300 400 1 1
197 40 0 36781 9000 400 400 1 1
198 41 0 14172 9000 3400 400 1 1
199 42 0 38982 9000 4200 400 1 1
200 43 0 144924 9000 20400 400 1 1
201 44 0 79250 9000 2400 400 1 1
202 45 0 135890 9000 44400 400 1 1
203 46 0 1071 9000 7700 400 1 1
204 47 0 476 9000 200 400 1 1
205 48 0 63914 9000 1300 400 1 1

```

```

206 49 0 2568 9000 600 400 1 1
207 50 0 21305 9000 3300 400 1 1
208 51 0 266687 9000 10290 400 1 1
209 52 0 78462 9000 4900 400 1 1
210 53 0 1060 9000 500 400 1 1
211 54 0 123 9000 200 400 1 1
212 55 0 55 9000 100 400 1 1
213 56 0 471 9000 200 400 1 1
214 57 0 16370 9000 200 400 1 1
215 58 0 140310 9000 1400 400 1 1
216 59 0 9154 9000 5400 400 1 1
217 }
218
219
220 @ASIC 0 {
221 # price StPwr freq pins BEvSleep BEvIDLE Area DVS
222 400 10 2500 21 0.01 512 18374 0
223
224 # Vmax vt CommBuffer CommTime CommPower CommArea
225 2.12196 0.402969 1 0.124688 21 32
226
227 #-----
228 # type ExeCyc DynPwr Area Exable
229 0 222 220 1233 1
230 1 142 141 2320 1
231 2 40 70 438 1
232 3 125 324 1549 1
233 4 146 460 2623 1
234 5 42 60 93 1
235 6 232 824 4321 0
236 7 120 390 1924 1
237 8 32 75 240 0
238 9 41 80 420 1
239 10 652 360 2324 1
240 11 868 376 2364 0
241 12 293 540 3522 0
242 13 221 540 3523 0
243 14 178 100 563 1
244 15 765 140 743 1
245 16 979 530 3983 1
246 17 132 150 865 1
247 18 142 330 1206 1
248 19 116 540 4432 1
249 20 128 270 1987 0
250 21 122 310 162 1
251 22 453 1000 8837 1
252 23 192 180 982 1
253 24 759 320 1634 1
254 25 267 420 1607 1
255 26 329 440 2230 1
256 27 114 110 541 0
257 28 196 200 733 1
258 29 78 260 1324 1
259 30 47 90 421 1
260 31 120 200 1065 1
261 32 136 190 836 0
262 33 115 300 1532 0
263 34 220 340 1728 1
264 35 643 670 3872 1
265 36 1270 540 2722 1
266 37 1809 920 5542 1
267 38 1980 870 6234 1
268 39 1220 420 3233 1
269 40 620 220 1432 1
270 41 236 550 2989 1
271 42 971 860 5827 1
272 43 2371 440 2251 1
273 44 1448 1210 10873 1
274 45 2112 1160 7691 1
275 46 582 630 2344 1

```

```

276 47 146 160 574 1
277 48 1282 790 3515 1
278 49 4324 870 5224 0
279 50 639 440 1734 1
280 51 5439 1720 12983 1
281 52 862 780 320 1
282 53 60 200 233 1
283 54 113 410 1092 1
284 55 155 520 2872 1
285 56 72 440 891 0
286 57 175 890 1541 1
287 58 1690 1920 2360 1
288 59 203 1030 832 1
289 }
290
291
292
293 @ASIC 1 {
294 # price StPwr freq pins BEVSleep BEVIDLE Area DVS
295 300 18 2500 24 0.029 2200 50000 0
296
297 # Vmax vt CommBuffer CommTime CommPower CommArea
298 3.3 1.8 1 0.35 37 52
299
300
301 #-----
302 # type ExeCyc DynPwr Area Exable
303 0 223 230 1244 1
304 1 162 70 2543 1
305 2 73 90 542 1
306 3 228 160 2675 1
307 4 148 160 1276 1
308 5 63 70 129 1
309 6 345 620 4899 0
310 7 178 270 1647 1
311 8 72 60 356 0
312 9 63 90 451 1
313 10 867 450 1345 1
314 11 1438 520 2452 1
315 12 328 410 3798 1
316 13 123 580 3748 0
317 14 438 80 623 1
318 15 542 70 719 1
319 16 1265 50 5098 0
320 17 245 130 956 0
321 18 356 150 1443 1
322 19 191 530 4932 1
323 20 342 140 1541 0
324 21 121 280 162 1
325 22 527 980 7476 1
326 23 267 150 1093 0
327 24 728 270 1897 1
328 25 189 380 1452 1
329 26 561 420 2786 1
330 27 544 90 618 1
331 28 259 200 781 1
332 29 87 250 1562 1
333 30 163 80 711 1
334 31 353 170 1379 1
335 32 178 190 855 1
336 33 190 240 1723 1
337 34 182 180 1615 1
338 35 1872 440 4412 1
339 36 1152 410 2437 1
340 37 1982 730 5994 1
341 38 2621 670 10841 1
342 39 1792 370 3257 1
343 40 901 180 2148 1
344 41 515 370 1756 1
345 42 874 770 5234 1

```

```

346 43 2538      410      2666      1
347 44 1983      1330     14134      1
348 45 3871      870      4111      1
349 46 773       720      2785      0
350 47 541       150      522       1
351 48 1826      660      3996      1
352 49 3992      740      5153      0
353 50 501       390      2214      1
354 51 2176     1210     13543      1
355 52 834       920      542       1
356 53 87        630      567       1
357 54 134       630      2398      1
358 55 182      1520     2001      1
359 56 98        860      562       0
360 57 289      880      1982      1
361 58 1469     2900     2120      1
362 59 323      2210     863       1
363 }
364
365
366 @FPGA 0 {
367 # price      StPwr      freq      pins      BEvSleep  BEvIDLE  CLBs      DVS
368 250          127        2500     53        0.058    1759    1220      0
369
370 # Vmax       vt          CommBuffer CommTime  CommPower  CommCLB  ReconT    ReconPwr
371 2.5          0.74        1        0.4339    42.79    213     1000      25.4241
372
373 #-----
374 # type ExeCyc      DynPwr      CLB      Exable
375 0      288        1820     52       0
376 1      160        2590     74       0
377 2      120        910      26       0
378 3      163        2620     75       1
379 4       89        4480    128       1
380 5      298        390      11       1
381 6      240        8820    252       0
382 7      120        4410    126       1
383 8      153        810      23       0
384 9      152        1020     29       1
385 10     4547       11650   333       1
386 11     1519       5390    154       0
387 12      288       6050    173       0
388 13      361       6970    199       0
389 14     2362      1120     32       1
390 15     8586      1470     42       1
391 16      346      7300    211       0
392 17      172      1800     53       0
393 18      115      260      76       1
394 19       65      930     268       1
395 20       67      340     99       0
396 21      193      560     16       1
397 22     2615     16660   476       1
398 23      661      2000     57       0
399 24     2600      2980     85       1
400 25      601      2700     77       1
401 26      701      4800    137       1
402 27      314      1100     34       0
403 28      296      2660     76       0
404 29      143      5220    149       0
405 30      222      1960     56       0
406 31       98      2730     78       0
407 32      361      1820     52       0
408 33      250      2560     73       1
409 34      220      3110     89       1
410 35     4158      6550    187       1
411 36    13866      5390    154       1
412 37    18024      9660    276       1
413 38    17747    11480    328       1
414 39    12203      5740    164       1
415 40     3678      2350     67       1

```



```

416 41 1417 5110 146 1
417 42 3898 11240 321 1
418 43 14492 3890 111 1
419 44 7925 18620 532 1
420 45 13589 12780 365 1
421 46 407 4480 128 0
422 47 276 1190 34 1
423 48 9178 6020 172 1
424 49 25755 9280 265 0
425 50 6391 2870 82 1
426 51 22668 23730 678 1
427 52 1853 920 542 1
428 53 687 4630 567 1
429 54 1023 7630 398 0
430 55 312 19520 402 1
431 56 1091 5760 162 0
432 57 932 9830 282 1
433 58 21469 22900 420 0
434 59 1323 2210 363 1
435 }
436
437
438
439 @COMM_AMOUNT 0 {
440 # type comamount
441 0 16
442 1 330
443 2 320
444 3 36
445 4 16
446 5 42
447 6 44
448 7 256
449 8 26
450 9 28
451 10 240
452 11 80
453 12 324
454 13 100
455 14 160
456 15 10
457 16 2
458 17 32
459 }
460
461
462 @LINK 0 {
463 # price StPwr freq pins BEvSleep maxUser PckSize PckOver
464 65 775 66000 15 0.3 7 8 33
465
466 # wTimeRate aPower
467 17846.3 4881.648
468 }
469
470 @LINK 1 {
471 # price StPwr freq pins BEvSleep maxUser PckSize PckOver
472 48 941 33000 15 0.5 7 8 33
473
474 # wTimeRate aPower
475 17846.3 4781.648
476 }
477
478
479 @MEM 0 {
480 # price StPwr size IOPwr
481 17.2338 89.7391 257586 145.237
482 }

```

Appendix C

DVS Problem Formulation for Distributed Heterogeneous Systems

As mentioned in Chapter 3, previous DVS approaches [19, 64, 89] identify scaling voltages without consideration of the power variations, unlike the PV-DVS optimisation presented here. The problem can be stated as follows:

Find for all DVS-PE mapped tasks $\tau \in \mathcal{T}_{\text{DVS}}$ of the system specification a single scaling voltage V_{dd} (between the threshold voltage V_t and the nominal supply voltage V_{max}) under consideration of individual power dissipations $P_{max}(\tau)$ such that the dynamic energy dissipation E_{Σ} is minimised and no deadline and precedence constraints are violated.

The problem can be mathematically expressed using the following definitions, where $\mathbb{R}_0^+ = \{x | x \in \mathbb{R}, 0 \leq x < +\infty\}$ and $\mathbb{R}^+ = \mathbb{R}_0^+ \setminus 0$:

- $G_S(\mathcal{T}, \mathcal{C})$ is the system specification graph, where \mathcal{T} is the set of tasks and \mathcal{C} is the set of communications.
- $G_A(\mathcal{P}, \mathcal{L})$ is a directed architecture graph, where \mathcal{P} is the set of PEs and \mathcal{L} is the set of CLs.
- $\mathcal{A} = \mathcal{T} \cup \mathcal{C}$ defines the set of all activities
- $\mathcal{K} = \mathcal{P} \cup \mathcal{L}$ defines the set of all allocated components

- $\mathcal{T}_{DVS} \subseteq \mathcal{T}$ denotes the set of all tasks mapped to DVS-PEs
- $P_{max} : \mathcal{T} \mapsto \mathbb{R}^+$ is a function returning the power dissipation of task τ executed at maximal PE supply voltage V_{max}
- $t_{min} : \mathcal{T} \mapsto \mathbb{R}^+$ is a function returning the minimal execution time of task $\tau \in \mathcal{T}$ at maximal PE supply voltage V_{max}
- $V_t : \mathcal{T} \mapsto \mathbb{R}^+$ is defined as a function which returns the threshold voltage of the PE to which task $\tau \in \mathcal{T}$ is mapped
- $V_{max} : \mathcal{T} \mapsto \mathbb{R}^+$ is a function returning the maximal supply voltage of the PE to which task $\tau \in \mathcal{T}$ is mapped
- $\mathcal{T}_d \subseteq \mathcal{T}$ denotes the set of all tasks having a hard deadline
- $t_{exe} : \mathcal{A} \mapsto \mathbb{R}^+$ is a function defined by:

$$t_{exe} = \begin{cases} t_{min}(\epsilon) \cdot \frac{V_{dd}(\epsilon)}{(V_{dd}(\epsilon) - V_t(\epsilon))^2} \cdot \frac{(V_{max}(\epsilon) - V_t(\epsilon))^2}{V_{max}(\epsilon)} & \text{if } \epsilon \in \mathcal{T} \\ t_C & \text{if } \epsilon \in \mathcal{C} \end{cases}$$

where t_C is the communication time for the communication activity $\gamma \in \mathcal{C}$

- $t_d : \mathcal{T}_d \mapsto \mathbb{R}_0^+$ is a function returning the deadline of task $\tau \in \mathcal{T}_d$
- $\mathcal{C}^{in} : \mathcal{T} \mapsto 2^{\mathcal{C}}$ returns the set of all ingoing edges of task $\tau \in \mathcal{T}$
- $t_S : \mathcal{E} \mapsto \mathbb{R}_0^+$ is a function which returns the start time of an activity $\epsilon \in \mathcal{A}$
- $\mathbf{A} : \mathcal{K} \mapsto 2^{\mathcal{A}}$ defines a function, returning the set of all activities mapped to a component $\kappa \in \mathcal{K}$
- $I = [t_S(\epsilon), (t_S(\epsilon) + t_{exe}(\epsilon))]$ is the execution interval of activity $\epsilon \in \mathcal{A}$
- $\mathbf{i} : \mathcal{A} \mapsto 2^I$ is a function returning the execution interval of an activity $\epsilon \in \mathcal{A}$

Using this definitions it is possible to formalise the problem mathematically as the minimisation of

$$E_{\Sigma} = \sum_{\tau \in \mathcal{T}_{DVS}} \frac{P_{max}(\tau) \cdot t_{min}(\tau)}{V_{max}^2(\tau)} \cdot V_{dd}^2(\tau) \quad (\text{C.1})$$

subject to

$$V_t(\tau) < V_{dd}(\tau) \leq V_{max}(\tau), \forall \tau \in \mathcal{T}_{DVS} \quad (\text{C.2})$$

$$\mathbf{t}_S(\tau) + \mathbf{t}_{exe}(\tau) \leq \mathbf{t}_d(\tau), \forall \tau \in \mathcal{T}_d \quad (\text{C.3})$$

$$\mathbf{t}_S(\gamma) + \mathbf{t}_{exe}(\gamma) \leq \mathbf{t}_S(\tau), \forall \tau \in \mathcal{T}, \gamma \in \mathbf{C}^{in}(\tau) \quad (\text{C.4})$$

$$\mathbf{i}(\varepsilon_n) \cap \mathbf{i}(\varepsilon_m) = \emptyset, \forall (\varepsilon_n, \varepsilon_m) \text{ so that } \varepsilon_n \in \mathbf{A}(\kappa_1), \varepsilon_m \in \mathbf{A}(\kappa_2) \Rightarrow \kappa_1 = \kappa_2 \quad (\text{C.5})$$

Appendix D

Derivation of the Energy/Delay Trade-Off

This appendix outlines the derivation of the energy/delay trade-off equation introduced in Chapter 2. The normalised circuit delay d^* is given by the increased circuit delay at $d(V_{dd})$ at reduced supply voltage V_{dd} over the minimal circuit delay $d(V_{max})$ at maximal supply voltage V_{max} . Using Equation (2.8) (Section 2.1, Page 26) the normalised delay is expressed as:

$$d^* = \frac{d(V_{dd})}{d(V_{max})} = \frac{V_{dd}}{(V_{dd} - V_t)^2} \cdot \frac{(V_{max} - V_t)^2}{V_{max}} \quad (D.1)$$

substituting

$$V_0 = \frac{(V_{max} - V_t)^2}{V_{max}} \quad (D.2)$$

results in

$$\frac{d^*}{V_0} = \frac{V_{dd}}{(V_{dd} - V_t)^2} \quad (D.3)$$

This can be transformed into the quadratic polynomial:

$$V_{dd}^2 - \left(2V_t + \frac{V_0}{d^*}\right) \cdot V_{dd} + V_t^2 = 0 \quad (D.4)$$

Solving the quadratic equation results in:

$$V_{dd} = V_t + \frac{V_0}{2d^*} + \sqrt{\left(V_t + \frac{V_0}{2d^*}\right)^2 - V_t^2} \quad (D.5)$$

From Equation (2.6), the normalised energy dissipation is given as:

$$\frac{E(V_{dd})}{E(V_{max})} = \frac{N_C \cdot \alpha \cdot C_L \cdot V_{dd}^2}{N_C \cdot \alpha \cdot C_L \cdot V_{max}^2} = \frac{V_{dd}^2}{V_{max}^2} \quad (D.6)$$

Inserting Equation (D.5) into Equation (D.6) results in the energy/delay trade-off equation:

$$E^*(d^*) = \frac{E(V_{max})}{V_{max}^2} \cdot \left(V_t + \frac{V_0}{2d^*} + \sqrt{\left(V_t + \frac{V_0}{2d^*} \right)^2 - V_t} \right)^2 \quad (D.7)$$

Note, the energy/delay trade-off depends on the energy dissipation of each task, i.e., on the individual power dissipation over time. Therefore, using this equation the different power dissipation on a single processing element are taken into account.

References

- [1] GNU CC Manual.
available at: <http://gcc.gnu.org/>.
- [2] GNU GCJ: GNU Compiler for Java.
available at: <http://gcc.gnu.org/java/>.
- [3] GNU gprof Manual.
available at: <http://www.gnu.org/manual/gprof-2.9.1/gprof.html>.
- [4] GSM 06.10, Technical University of Berlin.
Source code available at <http://kbs.cs.tu-berlin.de/~jutta/toast.html>.
- [5] Independent JPEG Group: jpeg-6b.
Source code available at <ftp://ftp.uu.net/graphics/jpeg/jpegsrc.v6b.tar.gz>.
- [6] International Technology Roadmap for Semiconductors 2000.
<http://public.itrs.net>.
- [7] Standard Performance Evaluation Corporation.
<http://www.specbench.org/>.
- [8] Synopsys Behavioral Compiler.
http://www.synopsys.com/products/beh_syn/beh_syn.html.
- [9] Synopsys Design Compiler.
http://www.synopsys.com/products/logic/design_comp_ds.html.
- [10] Synplify from Synplicity.
<http://www.synplicity.com/products/>.
- [11] The Free-IP Project.
<http://www.free-ip.com/>.
- [12] WITAS: The Wallenberg laboratory for research on Information Technology and Autonomous System.
<http://www.ida.liu.se/ext/witas/>.
- [13] Intel® XScale™ Core, Developer's Manual, December 2000. Order Number 273473-001.

- [14] Mobile AMD Athlon™4, Processor Model 6 CPGA Data Sheet, November 2000. Publication No 24319 Rev E.
- [15] Intel® PXA800F Cellular Processor, Developer's Manual, February 2003. Order Number 252569-002.
- [16] T. Adam, K. Chandy, and J. Dickson. A Comparison of List Scheduling for Parallel Processing Systems. *J. Communications of the ACM*, 17(12):685–690, December 1974.
- [17] Thomas Bäck, Ulrich Hammel, and Hans-Paul Schwefel. Evolutionary Computation: Comments on the History and Current State. *IEEE Transactions on Evolutionary Computation*, 1(1):3–17, 1997.
- [18] Z. Baidas, A. D. Brown, and A. C. Williams. Floating-point Behavioral Synthesis. *IEEE Transactions on Computer-Aided Design*, 20(7):828–839, July 2001.
- [19] N. Bambha, S. Bhattacharyya, J. Teich, and E. Zitzler. Hybrid Global/Local Search Strategies for Dynamic Voltage Scaling in Embedded Multiprocessors. In *Proceedings 1st International Symposium Hardware/Software Co-Design (CODES'01)*, pages 243–248, April 2001.
- [20] Armin Bender. Design of an Optimal Loosely Coupled Heterogeneous Multiprocessor System. In *Proceedings European Design Automation Conference*, pages 275–281, March 1996.
- [21] L. Benini, G. De Micheli, E. Macii, D. Sciuto, and C. Silvano. Address Bus Encoding Techniques for System-Level Power Optimization. In *Proceedings Design, Automation and Test in Europe Conference (DATE98)*, pages 861–866, March 1998.
- [22] Luca Benini, A. Bogliolo, and Giovanni De Micheli. A Survey of Design Techniques for System-Level Dynamic Power Management. *IEEE Transactions on VLSI Systems*, pages 299–316, June 2000.
- [23] Luca Benini, A. Bogliolo, G.A. Paleologo, and G. De Micheli. Policy Optimization for Dynamic Power Management. *IEEE Transactions on Computer-Aided Design*, 18(6):813–833, June 1999.
- [24] Luca Benini, Alessandro Bogliolo, and Giovanni De Mecheli. Dynamic Power Management of Electronic Systems. In *Proceedings IEEE/ACM International Conference Computer-Aided Design (ICCAD-98)*, pages 696–702, Nov 1998.
- [25] Luca Benini and Giovanni De Micheli. *Dynamic Power Management: Design Techniques and CAD Tools*. Kluwer Academic Publishers, 1997.
- [26] Luca Benini, Giovanni De Micheli, Enrico Macii, Massimo Poncino, and R. Scarsi. Symbolic Synthesis of Clock-gating Logic for Power Optimization of Synchronous Controllers. *ACM Transactions on Design Automation of Electronic Systems*, 4(4):351–375, 1999.

- [27] Luca Benini, Polly Siegel, and Giovanni De Micheli. Saving Power by Synthesizing Gated Clocks for Sequential Circuits. *IEEE Design & Test of Computers*, 11(4):32–41, 1994.
- [28] Peter Bjørn-Jørgensen and Jan Madsen. Critical Path Driven Cosynthesis for Heterogeneous Target Architectures. In *Proceedings 5th International Workshop Hardware/Software Co-Design (Codes/CASHE'97)*, pages 15 – 19, 1997.
- [29] C. Brandolese, W. Fornaciari, F. Salice, and D. Sciuto. Energy Estimation for 32 bit Microprocessors. In *Proceedings 8th International Workshop Hardware/Software Co-Design (CODES'00)*, pages 24–28, May 2000.
- [30] Jason J. Brown, Danny Z.Chen, Garrison W. Greenwood, Xiaobo (Sharon) Hu, and Richard W. Taylor. Scheduling for Power Reduction in a Real-Time System. In *Proceedings International Symposium Low Power Electronics and Design (ISLPED'97)*, pages 84–87, 1997.
- [31] Thomas D. Burd. *Energy-Efficient Processor System Design*. PhD thesis, University of California at Berkeley, 2001.
- [32] Thomas D. Burd and Robert W. Brodersen. Processor Design for Portable Systems. *Journal on VLSI Signal Processing*, 13(2):203–222, August 1996.
- [33] Thomas D. Burd, Trevor A. Pering, Anthony J. Stratakos, and Robert W. Brodersen. A Dynamic Voltage Scaled Microprocessor System. *IEEE Journal on Solid-State Circuits*, 35(11):1571–1580, November 2000.
- [34] Anantha P. Chandrakasan and Robert W. Brodersen. *Low Power Digital CMOS Design*. Kluwer Academic Publisher, 1995.
- [35] Anantha P. Chandrakasan, T. Sheng, and Robert W. Brodersen. Low Power CMOS Digital Design. *Journal of Solid State Circuits*, 27(4):473–484, April 1992.
- [36] J. D'Ambrosio and X. Hu. Configuration-Level Hardware/Software Partitioning for Real-Time Embedded Systems. In *Proceedings International Workshop Hardware/Software Co-Design (Codes/CASHE'94)*, pages 34–41, 1994.
- [37] Bharat P. Dave, Ganesh Lakshminarayana, and Niraj K. Jha. COSYN: Hardware-Software Co-Synthesis of Embedded Systems. In *Proceedings IEEE 34th Design Automation Conference (DAC97)*, pages 703–708, 1997.
- [38] G. DeMicheli. *Synthesis and Optimization of Digital Circuits*. McGraw-Hill, 1994.
- [39] Giovanni DeMicheli, David C. Ku, Frederic Mailhot, and Thomas Truong. The Olympus Synthesis System for Digital Design. *IEEE Design & Test of Computers*, pages 37–53, October 1990.
- [40] Micheal L. Dertouzos and Aloysius Ka-Lau Mok. Multiprocessor On-Line Scheduling of Hard-Real-Time Tasks. *IEEE Transactions on Software Engineering*, 15(12):1497–1506, December 1989.

- [41] Srinivas Devadas and Sharad Malik. A Survey of Optimization Techniques Targeting Low Power VLSI Circuits. In *Proceedings IEEE 32nd Design Automation Conference (DAC95)*, pages 242–247, 1995.
- [42] Muhammad K. Dhodhi, Imtiaz Ahmad, and Robert Storer. SHEMUS: Synthesis of Heterogeneous Multiprocessor Systems. *J. Microprocessors and Microsystems*, 19(6):311–319, August 1995.
- [43] R. Dick and N. K. Jha. MOCSYN: Multiobjective core-based single-chip system synthesis. In *Proceedings Design, Automation and Test in Europe Conference (DATE99)*, pages 263–270, March 1999.
- [44] R. Dick, D. Rhodes, and W. Wolf. TGFF: Task Graphs for free. In *Proceedings 5th International Workshop Hardware/Software Co-Design (Codes/CASHE'97)*, pages 97–101, March 1998.
- [45] Robert P. Dick and Niraj K. Jha. MOGAC: A Multiobjective Genetic Algorithm for Hardware-Software Co-Synthesis of Distributed Embedded Systems. *IEEE Transactions on Computer-Aided Design*, 17(10):920–935, Oct 1998.
- [46] Petru Eles, Alexa Doboli, Paul Pop, and Zebo Peng. Scheduling with Bus Access Optimization for Distributed Embedded Systems. *IEEE Transactions on VLSI Systems*, 8(5):472–491, Oct 2000.
- [47] Petru Eles, Krzysztof Kuchcinski, Zebo Peng, Alexa Doboli, and Paul Pop. Scheduling of Conditional Process Graphs for the Synthesis of Embedded Systems. In *Proceedings Design, Automation and Test in Europe Conference (DATE98)*, pages 132–138, 1998.
- [48] Petru Eles, Zebo Peng, Krzysztof Kuchcinski, and Alexa Doboli. System Level Hardware/Software Partitioning Based on Simulated Annealing and Tabu Search. *Kluwer Journal on Design Automation for Embedded Systems*, 2:5–32, 1997.
- [49] R. Ernst and J. Henkel. Hardware-Software Codesign of Embedded Controllers Based on Hardware Extraction. In *1st International Workshop Hardware/Software Co-Design (Codes/CASHE'92)*, 1992.
- [50] R. Ernst, J. Henkel, and Th. Brenner. Hardware-Software Co-synthesis for Micro-Controllers. *IEEE Design & Test of Computers*, 10(4):64–75, Dec 1993.
- [51] Rolf Ernst. Codesign of Embedded Systems: Status and Trends. *IEEE Design & Test of Computers*, pages 45–54, April 1998.
- [52] A. Feller. Automatic layout of low-cost quick-turnaround random-logic custom LSI devices. In *Proceedings IEEE 13th Design Automation Conference (DAC76)*, pages 97–85, 1976.
- [53] Terence C. Fogarty. Varying the probability of mutation in the genetic algorithm. In *Proceedings 3rd International Conference on Genetic Algorithms (ICGA)*, pages 104–109, 1989.

- [54] W. Fornaciari, D. Sciuto, and C. Silvano. Power Estimation for Architectural Exploration of HW/SW Communication on System-Level Buses. In *Proceedings 7th International Workshop Hardware/Software Co-Design (CODES'99)*, pages 152–156, May 1999.
- [55] Michael L. Fredman and Robert Endre Tarjan. Fibonacci Heaps and Their Uses in Improved Network Optimization Algorithms. In *Annual Symposium on Foundations of Computer Science (FOCS 1984)*, 1984.
- [56] Daniel Gajski and Longanath Ramachandran. Introduction to High-Level Synthesis. *IEEE Design and Test of Computers*, 2(4):44–54, 1994.
- [57] Daniel D. Gajski, Jianwen Zhu, and Rainer Dömer. Essential Issues in Codesign. Technical report, University of California, Irvine, Department of Information and Computer Science, June 1997.
- [58] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the theory of NP-Completeness*. W.H. Freeman and Company, 1979.
- [59] Sabih H. Gerez. *Algorithm for VLSI Design Automation*. John Wiley & Sons Ltd., 1998.
- [60] David E. Goldberg. *Genetic Algorithms in Search, Optimization & Machine Learning*. Addison-Wesley Publishing Company, 1989.
- [61] James Goodman, Anantha Chandrakasan, and Abram P. Dancy. Design and Implementation of a Scalable Encryption Procesoor with Embedded Variable DC/DC Converter. In *Proceedings IEEE 36th Design Automation Conference (DAC99)*, pages 855–860, 1999.
- [62] Martin Grajcar. Genetic List Scheduling Algorithm for Scheduling and Allocation on a Loosely Coupled Heterogeneous Multiprocessor System. In *Proceedings IEEE 36th Design Automation Conference (DAC99)*, pages 280–285, 1999.
- [63] Flavius Gruian. System-Level Design Methods for Low-Energy Architectures Containing Variable Voltage Processors. In *Workshop Power-Aware Computing Systems*, Nov 2000.
- [64] Flavius Gruian and Krzysztof Kuchcinski. LEneS: Task Scheduling for Low-Energy Systems Using Variable Supply Voltage Processors. In *Proceedings Asia South Pacific - Design Automation Conference (ASP-DAC'01)*, pages 449–455, Jan 2001.
- [65] R. K. Gupta and G. De Micheli. Hardware/Software Co-synthesis of Digital Systems. *IEEE Design & Test of Computers*, pages 29–41, September 1993.
- [66] R.K. Gupta. *Co-Synthesis of Hardware and Software for Digital Embedded Systems*. PhD thesis, Stanford University, December 1993.
- [67] Vadim Gutnik and Anantha Chandrakasan. Embedded Power Supply for Low-Power DSP. *IEEE Transactions on VLSI Systems*, 5(4), 425–435 1997.

- [68] Johan Hagman. mpeg3play-0.9.6.
Source code available at <http://home.swipnet.se/~w-10694/tars/mpeg3play-0.9.6-x86.tar.gz>.
- [69] Lajos Hanzo, Clare Somerville, and Jason Woodard. *Voice Compression and Communications: Principles and Applications for Fixed and Wireless Channels*. John Wiley & Sons Inc., 2001.
- [70] Jörg Henkel. A Low Power Hardware/Software Partitioning Approach for Core-Based Embedded Systems. In *Proceedings IEEE 36th Design Automation Conference (DAC99)*, pages 122–127, 1999.
- [71] Jörg Henkel and Rolf Ernst. An Approach to Automated Hardware/Software Partitioning using a Flexible Granularity that is driven by High-Level Estimation Techniques. *IEEE Transactions on VLSI Systems*, 9(2):273–289, 2001.
- [72] Inki Hong, Darko Kirovski, Gang Qu, Miodrag Potkonjak, and Mani B. Srivastava. Power Optimization of Variable-Voltage Core-Based Systems. *IEEE Transactions on Computer-Aided Design*, 18(12):1702–1714, Dec 1999.
- [73] Junwei Hou and Wayne Wolf. Process Partitioning for Distributed Embedded Systems. In *Proceedings 4th International Workshop Hardware/Software Co-Design (Codes/CASHE'96)*, pages 70 – 76, March 1996.
- [74] Tohru Ishihara and Hiroto Yasuura. Voltage Scheduling Problem for Dynamically Variable Voltage Processors. In *Proceedings International Symposium Low Power Electronics and Design (ISLPED'98)*, pages 197–202, 1998.
- [75] Asawaree Kalavade and Edward A. Lee. A Global Criticality/Local Phase Driven Algorithm for the Constrained Hardware/Software Partitioning Problem. In *Proceedings International Workshop Hardware/Software Co-Design (Codes/CASHE'94)*, pages 42–48, Sept. 1994.
- [76] Asawaree Kalavade and P. A. Subrahmanyam. Hardware/Software Partitioning for Multifunction Systems. *IEEE Transactions on Computer-Aided Design*, 17(9):819–836, Sep 1998.
- [77] Kurt Keutzer, Sharad Malik, A. Richard Newton, Jan M. Rabaey, and A. Sangiovanni-Vincentelli. System-Level Design: Orthogonalization of Concerns and Platform-Based Design. *IEEE Transactions on Computer-Aided Design*, 19(12):1523–1543, December 2000.
- [78] Alexander Klaiber. The Technology behind Crusoe Processors. January 2000.
<http://www.transmeta.com>.
- [79] P. V. Knudsen and J. Madsen. Integrating Communication Protocol Selection with Hardware/Software Codesign. *IEEE Transactions on Computer-Aided Design*, 18(9):1077–1095, Aug 1999.

- [80] Peter V. Knudsen and Jan Madsen. PACE: A Dynamic Programming Algorithm for Hardware/Software Partitioning. In *Proceedings 4th International Workshop Hardware/Software Co-Design (Codes/CASHE'96)*, pages 85 – 92, March 1996.
- [81] Yu-Kwong Kwok and Ishfaq Ahmad. Dynamic Critical-Path Scheduling: An Effectiveness Technique for Allocating Task Graphs to Multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 7(5):506–521, May 1996.
- [82] Yu-Kwong Kwok and Ishfaq Ahmad. Static Scheduling Algorithms for Allocating Directed Task Graphs to Multiprocessors. *ACM Computing Surveys*, 31(4):406–471, December 1999.
- [83] Rainer Leupers and Peter Marwedel. *Retargetable Compiler Technology for Embedded Systems - Tools and Applications*. Kluwer Academic Publishers, 2001.
- [84] Y. Li and J. Henkel. A Framework for Estimating and Minimizing Energy Dissipation of Embedded HW/SW Systems. In *Proceedings IEEE 35th Design Automation Conference (DAC98)*, pages 188–193, 1998.
- [85] Y.-T. S. Li, S. Malik, and A. Wolfe. Performance Estimation of Embedded Software with Instruction Cache Modeling. In *Proceedings IEEE/ACM International Conference Computer-Aided Design (ICCAD-95)*, pages 380–387, November 1995.
- [86] Jinfeng Liu, Pai H. Chou, and Nader Bagherzadeh. Communication Speed Selection for Embedded Systems with Networked Voltage-Scalable Processors. In *Proceedings 2nd International Symposium Hardware/Software Co-Design (CODES'02)*, pages 169–174, 2002.
- [87] J. R. Lorch and A. J. Smith. Software Strategies for Portable Computer Energy Management. *IEEE Personal Communications*, 5(3):60–73, June 1998.
- [88] Yung-Hsiang Lu, Luca Benini, and Giovanni De Micheli. Low-Power Task Scheduling for Multiple Devices. In *Proceedings 8th International Workshop Hardware/Software Co-Design (CODES'00)*, pages 39–43, 2000.
- [89] Jiong Luo and Niraj K. Jha. Power-conscious Joint Scheduling of Periodic Task Graphs and Aperiodic Tasks in Distributed Real-time Embedded Systems. In *Proceedings IEEE/ACM International Conference Computer-Aided Design (ICCAD-00)*, pages 357–364, Nov 2000.
- [90] Jiong Luo and Niraj K. Jha. Battery-aware Static Scheduling for Distributed Real-Time Embedded Systems. In *Proceedings IEEE 38th Design Automation Conference (DAC01)*, pages 444–449, 2001.
- [91] Jan Madsen and Peter Bjørn-Jørgensen. Embedded System Synthesis under Memory Constraints. In *Proceedings 7th International Workshop Hardware/Software Co-Design (CODES'99)*, pages 188 – 192, 1999.
- [92] Sven Mattisson. Minimizing Power Dissipation of Cellular Phones. In *ISLPED*, pages 42–45, 1997.

- [93] G. De Micheli and R. K. Gupta. Hardware/Software Co-Design. In *Proceedings of the IEEE*, pages 349–365, March 1997.
- [94] Giovanni De Micheli, Rolf Ernst, and Wayne Wolf. *Readings in Hardware/Software Co-Design*. Morgan Kaufmann Publishers, 2002.
- [95] J. Monteiro, S. Devadas, and A. Ghosh. Retiming Sequential Circuits for Low Power. In *Proceedings IEEE/ACM International Conference Computer-Aided Design (ICCAD-93)*, pages 398–402, 1993.
- [96] J. Monteiro, S. Devadas, and A. Ghosh. Sequential logic optimization for low power using input disabling precomputation architectures. *IEEE Transactions on Computer-Aided Design*, 17(3):279–284, March 1998.
- [97] W. Namgoong, M. Yu, and T. Meng. A High-Efficiency Variable-Voltage CMOS Dynamic dc-dc Switching Regulator. In *Proceedings International Solid-State Circuits Conference*, pages 380–381, 1997.
- [98] David Naylor and Simon Jones. *VHDL: A Logic Synthesis Approach*. Chapman-Hall, 1997.
- [99] Hyunok Oh and Soonhoi Ha. A Static Scheduling Heuristic for Heterogeneous Processors. In *2nd International EuroPar Conference Vol. II*, August 1996.
- [100] Hyunok Oh and Soonhoi Ha. A Hardware-Software Cosynthesis Technique Based on Heterogeneous Multiprocessor Scheduling. In *Proceedings 7th International Workshop Hardware/Software Co-Design (CODES'99)*, pages 183–187, May 1999.
- [101] Hyunok Oh and Soonhoi Ha. Hardware-Software Cosynthesis of Multi-Mode Multi-Task Embedded Systems with Real-Time Constraints. In *Proceedings 2nd International Symposium Hardware/Software Co-Design (CODES'02)*, pages 133–138, May 2002.
- [102] P. R. Panda and N. D. Dutt. Reducing Address Bus Transitions for Low Power Memory Mapping. In *IEEE European Design and Test Conference*, pages 63–67, March 1996.
- [103] Massoud Pedram. Power Minimization in IC Design: Principles and Applications. *ACM Transactions Design Automation of Electronic Systems (TODAES)*, 1(1):3–56, Jan 1996.
- [104] P. Pedro and Alan Burns. Schedulability Anaysis for Mode Changes in Flexible Real-time Systems. In *Proceddings Euromicro Workshop on Real-Time Systems*, pages 17–19, June 1998.
- [105] Paul Pop, Petru Eles, and Zebo Peng. Bus Access Optimization for Distributed Embedded Systems Based on Schedulability Analysis. In *Proceedings Design, Automation and Test in Europe Conference (DATE2000)*, pages 567–574, 2000.

- [106] Paul Pop, Petru Eles, and Zebo Peng. Scheduling with Optimized Communication for Time-Triggered Embedded Systems. In *Proceedings 8th International Workshop Hardware/Software Co-Design (CODES'00)*, pages 62–66, 2000.
- [107] Paul Pop, Petru Eles, Traian Pop, and Zebo Peng. An Approach to Incremental Design of Distributed Embedded Systems. In *Proceedings IEEE 38th Design Automation Conference (DAC01)*, pages 450–455, 2001.
- [108] S. Prakash and A. Parker. SOS: Synthesis of Application-Specific Heterogeneous Multiprocessor Systems. *J. Parallel & Distributed Computing*, pages 338–351, Dec 1992.
- [109] A. Raghunathan and N.K. Jha. SCALP: An iterative improvement based low-power data path synthesis algorithm. *IEEE Transactions on Computer-Aided Design*, 16(11):1260–1277, November 1997.
- [110] A. Raghunathan, N. K. Niraj, and Sujit Dey. *High-Level Power Analysis and Optimization*. Kluwer Academic Publishers, 1998.
- [111] Krithi Ramamritham and John A. Stankovic. Scheduling Algorithms and Operating Systems Support for Real-time Systems. *Proceedings of the IEEE*, 81(1):55–67, 1994.
- [112] R. L. Rhodes and Wayne Wolf. Co-Synthesis of Heterogeneous Multiprocessor Systems using Arbitrated Communication. In *Proceedings IEEE/ACM International Conference Computer-Aided Design (ICCAD-99)*, pages 339–342, 1999.
- [113] *International Technology Roadmap for Semiconductors*.
<http://notes.sematech.org/ntrs/PublNTRS.nsf>.
- [114] Alex Rogers and Adam Prügel-Bennett. Modelling the dynamics of a steady-state genetic algorithm. In *Foundations of Genetic Algorithms (FOGA-5)*, pages 57–68. Sept 1999.
- [115] James A. Rowson. Hardware/Software Co-Simulation. In *Proceedings IEEE 31st Design Automation Conference (DAC94)*, pages 439–440, 1994.
- [116] Kaushik Roy and Sharat C. Prasad. *Low-Power CMOS VLSI Circuit Design*. Wiley-Interscience, 2000.
- [117] Andrew Rushton. *VHDL for Logic Synthesis*. John Wiley & Son Ltd, 1998.
- [118] D. Saha, R. S. Mitra, and A. Basu. Hardware Software Partitioning using Genetic Algorithm. In *10th International Conference on VLSI Design*, pages 155–160, January 1997.
- [119] Marcus T. Schmitz and Bashir M. Al-Hashimi. Synthesising Low Power Distributed Systems: A Multi-Mode System Approach. Presented at the 2nd UK ACM SIGDA Workshop 2003, 16–17 September 2003.

- [120] Marcus T. Schmitz and Bashir M. Al-Hashimi. Energy Minimisation for Processor Cores using Variable Supply Voltages. In *Proceedings IEE Systems on a chip Workshop*, pages 101–104, Sept 2000.
- [121] Marcus T. Schmitz and Bashir M. Al-Hashimi. Low Power Process Assignment for Distributed Embedded Systems Using Dynamic Voltage Scaling. In *IEE Hardware/Software Codesign Workshop, London*, December 2000.
- [122] Marcus T. Schmitz and Bashir M. Al-Hashimi. Considering Power Variations of DVS Processing Elements for Energy Minimisation in Distributed Systems. In *Proceedings International Symposium System Synthesis (ISSS'01)*, pages 250–255, October 2001.
- [123] Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles. Energy-Efficient Mapping and Scheduling for DVS Enabled Distributed Embedded Systems. In *Proceedings Design, Automation and Test in Europe Conference (DATE2002)*, pages 514–521, March 2002.
- [124] Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles. Synthesizing Energy-Efficient Embedded Systems with LOPOCOS. *Kluwer Journal on Design Automation for Embedded Systems*, 6(4):401–424, July 2002.
- [125] Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles. A Co-Design Methodology for Energy-Efficient Multi-Mode Embedded Systems with Consideration of Mode Execution Probabilities. In *Proceedings Design, Automation and Test in Europe Conference (DATE2003)*, pages 960–965, March 2003.
- [126] Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles. Iterative Schedule Optimisation for Voltage Scalable Distributed Embedded Systems. *accepted for publication in ACM Transactions on Embedded System Design*, 2003.
- [127] Marcus T. Schmitz, Bashir M. Al-Hashimi, and Petru Eles. System-Level Design of Energy-Efficient Multi-Mode Embedded Systems. *currently under preparation for submission to IEEE Transactions on CAD*, 2003.
- [128] E. M. Sentovitch, K. J. Singh, Luciano Lavagno, Cho Moon, Rajeev Murgai, Alexander Saldanha, Hamid Savoj, Paul R. Stephan, Robert K. Brayton, and Alberto Sangiovanni-Vincentelli. SIS: A System for Sequential circuit synthesis. Technical report, University of California, Berkeley, May 1992.
- [129] L. Sha, R. Rajkumar, J. Lehoczky, and K Ramamritham. Mode Change Protocols for Priority-driven Preemptive Scheduling. 1:243–265, December 1989.
- [130] Y. Shin, D. Kim, and K. Choi. Schedulability-Driven Performance Analysis of Multiple Mode Embedded Real-Time Systems. In *Proceedings IEEE 37th Design Automation Conference (DAC00)*, pages 495–500, June 2000.
- [131] Youngsoo Shin and Kiyoun Choi. Power Conscious Fixed Priority Scheduling for Hard Real-Time Systems. In *Proceedings IEEE 36th Design Automation Conference (DAC99)*, pages 134–139, 1999.

- [132] Gilbert C. Sih and Edward A. Lee. A Compile-time scheduling heuristic for interconnection-constrained heterogeneous processor architectures. *IEEE Transactions on Parallel and Distributed Systems*, 4(2):175–187, February 1993.
- [133] T. Simunic, L. Benini, and G. De Micheli. Energy-Efficient Design of Battery-Powered Embedded Systems. In *Proceedings International Symposium Low Power Electronics and Design (ISLPED'99)*, pages 212–217, 1999.
- [134] M. Srivastava, A. Chandrakasan, and R. Brodersen. Predictive System Shutdown and other Architectural Techniques for Energy Efficient Programmable Computations. *IEEE Transactions on VLSI Systems*, 4(1):42–55, March 1996.
- [135] M. R. Stan and W. P. Burleson. Bus-Inverse Coding for Low-Power I/O. *IEEE Transactions on VLSI Systems*, 3(1):49–58, March 1995.
- [136] Jorgen Staunstrup and Wayne Wolf. *Hardware/Software Co-Design: Principles and Practice*. Kluwer Academic Publishers, 1997.
- [137] J. Teich, T. Blickle, and Lothar Thiele. An Evolutionary Approach to System-Level Synthesis. In *Proceedings 5th International Workshop Hardware/Software Co-Design (Codes/CASHE'97)*, pages 167 – 171, March 1997.
- [138] V. Tiwari and M. Lee. Power Analysis of a 32-bit Embedded Microcontroller. In *ASP-DAC*, pages 141–148, Aug 1995.
- [139] V. Tiwari, S. Malik, A. Wolfe, and M. Lee. Instruction Level Power Analysis and Optimization of Software. *Journal of VLSI Signal Processing Systems*, 13(2–3):223–238, 1996.
- [140] Vivek Tiwari, Sharad Malik, and Andrew Wolfe. Power Analysis of Embedded Software: A First Step Towards Software Power Minimization. *IEEE Transactions on VLSI Systems*, Dec 1994.
- [141] C-Y. Tsui, J. Monteiro, M. Pedram, S. Devadas, A. Despain, and B. Lin. Precomputation-Based Sequential Logic Optimization for Low Power. *IEEE Transactions on VLSI Systems*, 2(4):426–436, December 1994.
- [142] Bhaskaran Vasuder. *Image and Video Compression Standards*. Kluwer Academic Publisher, 1997.
- [143] Matthew Wall. GALib: A C++ Library of Genetic Algorithm Components Version 2.45, August 1996.
available at: <http://lancet.mit.edu/ga>.
- [144] M. Weiser, B. Welch, A. Demers, and S. Shenker. Scheduling for Reduced CPU Energy. In *Proceedings USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 13–23, 1994.
- [145] T. Wangtong, Peter Y.K. Cheung, and W. Luk. Comparing Three Heuristic Search Methods for Functional Partitioning in Hardware-Software Codesign. *Design Automation for Embedded Systems*, 6(4):425–449, July 2002.

- [146] A. C. Williams. *A Behavioural VHDL Synthesis System using Data Path Optimisation*. PhD thesis, University of Southampton, October 1997.
- [147] A. C. Williams, A. D. Brown, and M. Zwolinski. Simultaneous optimisation of dynamic power, area and delay in behavioural synthesis. *IEE Proc.-Comput. Digit. Tech.*, 147(6):383–390, November 2000.
- [148] Wayne H. Wolf. Hardware/Software Co-Design of Embedded Systems. In *Proceedings of the IEEE*, pages 967–989, July 1994.
- [149] Wayne H. Wolf. An Architectural Co-Synthesis Algorithm for Distributed, Embedded Computing Systems. *IEEE Transactions on VLSI Systems*, 5(2):218–229, June 1997.
- [150] M. Wu and D. Gajski. Hypertool: A Programming Aid for Message-passing Systems. *IEEE Transactions on Parallel and Distributed Systems*, 1(3):330–343, July 1990.
- [151] Y. Xie and Wayne Wolf. Allocation and Scheduling of Conditional Task Graph in Hardware/Software Co-Synthesis. In *Proceedings Design, Automation and Test in Europe Conference (DATE2001)*, pages 620 – 625, March 2001.
- [152] Frances Yao, Alan Demers, and Scott Shenker. A Scheduling Model for Reduced CPU Energy. In *IEEE Symposium on Foundations of Comp. Science*, pages 374–382, 1995.