

A VDM-RT METHODOLOGY
FOR THE
HW/SW CO-DESIGN
OF
EMBEDDED SYSTEMS

BY
JOSÉ ANTONIO ESPARZA ISASA
20097706

MASTER'S THESIS
IN
TECHNICAL INFORMATION TECHNOLOGY
DISTRIBUTED REAL-TIME SYSTEMS

SUPERVISOR: PROF. PETER GORM LARSEN
Aarhus School of Engineering
December 15, 2011

José Antonio Esparza Isasa
20097706

Peter Gorm Larsen
Supervisor



FACULTY OF SCIENCE
AARHUS UNIVERSITY

Abstract

This thesis presents a new methodology for Hardware/Software Co-design of embedded systems. The proposed methodology is based on the modelling language VDM-RT. This new approach allows system engineers to perform design space exploration at an abstract level earlier in the development process. This methodology supports the analysis of real-time deadlines under different hardware/software architectures, without making use of virtual or actual prototyping. Based on the information gained during this analysis, system engineers are able to allocate the required system functionality in hardware and software blocks. This allocation or partitioning process can be done through a rigorous study of the design trade-offs by applying the VDM-RT methodology. As an additional advantage, the system engineers are able to specify and capture clearly and unambiguously the system structure and behaviour. Besides presenting this new methodology, this thesis will show its application in two case studies and present a pragmatic analysis of its performance.

Acknowledgements

I would like to thank my academic supervisor Peter Gorm Larsen for his attention and advice during this thesis and my Master's degree. I would like to thank my industrial supervisors from Bang & Olufsen Gert Schilkowski and Karsten Langhoff Sørensen for their contribution in the AVB case study. Finally, I would like to thank Joey Coleman for his input regarding the content and structure of this work. Additional thanks go to my family and friends for their patience and support.

Table of Contents

Abstract	i
Acknowledgements	iii
Table of Contents	v
List of Figures	ix
Chapter 1 Introduction	1
1.1 Background	1
1.2 Motivation	2
1.3 Thesis goal, approach and scope	2
1.3.1 Goals	2
1.3.2 Approach	3
1.3.3 Scope	3
1.4 Case studies	4
1.5 Related work	4
1.6 Reading guide	6
1.7 Structure	7
Chapter 2 System development and abstraction in hardware/software engineering	9
2.1 Introduction	9
2.2 Abstraction levels overview in a single system	10
2.3 The System-Level perspective	12
2.3.1 Different approaches in system-level design	12
2.3.2 Application of system-level design approaches	12
2.4 Modelling languages in system-level design	13
2.5 Hardware/Software co-design	15
2.5.1 Trade-offs and impact of Hardware/Software co-design	15
2.5.2 Overview of the Hardware/Software co-design process	18
2.5.3 The role of modelling in Hardware/Software co-design	19
2.5.4 Outcomes of Hardware/Software co-design	20
2.5.5 Enabling technologies	21
2.6 Summary	22
Chapter 3 Modelling techniques, analysis and methodologies in hardware/software co-design	23
3.1 Introduction	23
3.2 Modelling languages	24

Table of Contents

3.2.1	SysML	24
3.2.2	SystemC	26
3.2.3	MATLAB - Simulink	27
3.2.4	POOSL	28
3.2.5	VDM-RT	29
3.2.6	Language comparison	31
3.3	Time and modelling languages	31
3.4	Methodologies	33
3.4.1	Xilinx methodology	34
3.4.2	A MATLAB - Simulink based methodology	34
3.4.3	A combined SysML - SystemC based methodology	35
3.4.4	The SHE methodology	35
3.4.5	A VDM-RT based methodology	37
3.5	Summary	38
Chapter 4 A VDM-RT based modelling methodology for Hardware/Software engineering		41
4.1	Introduction	41
4.2	Methodology description	41
4.2.1	Creation of a sequential model of the system	42
4.2.2	Creation of a concurrent model of the system	43
4.2.3	Creation of a distributed real-time model	43
4.2.4	Evaluating hardware partitions from VDM-RT models	44
4.3	Early prototype generation and model validation	47
4.4	Proposed improvements to the VDM technology	48
4.5	Modelling new platforms under VDM-RT	49
4.6	Summary	50
Chapter 5 Introductory case study: RC servo control		51
5.1	Introduction to basic servo control	51
5.2	Modelling of a servo controller	52
5.2.1	Sequential model	53
5.2.2	Concurrent model	55
5.2.3	Distributed real-time model	57
5.3	Implementation of the servo controller	63
5.3.1	Preliminary details	63
5.3.2	Implementing the conflictive architecture	63
5.3.3	Implementing the optimal architecture	64
5.3.4	RTOS based solution	66
5.4	Evaluating the implemented architectures	67
5.4.1	Evaluation of the conflictive architecture	67
5.4.2	Evaluation of the optimal architecture	69
5.4.3	Physical results	70
5.5	Summary	70
Chapter 6 Case study: Development of an AVB endpoint device		73
6.1	Introduction	73
6.2	Description of the AVB protocol	73
6.2.1	The traffic shaping protocol 802.1Qav	74

Table of Contents

6.2.2	The stream reservation protocol 802.1Qat	75
6.2.3	The time synchronization protocol 802.1As	75
6.3	Partitioning decisions and modelling the AVB compliant device	80
6.3.1	Protocol modelling	80
6.3.2	Real-time characteristics	81
6.3.3	Partitioning factors to be considered by the models	81
6.4	Modelling the AVB time synchronization protocol	81
6.4.1	Sequential model	81
6.4.2	Concurrent model	82
6.4.3	Distributed real-time models	83
6.5	Evaluating different hardware/software architectures in a simulated deployment	85
6.5.1	Architecture 1: highly software based architecture	85
6.5.2	Architecture 2: mostly software and offloaded SlavePort	86
6.5.3	Architecture 3: offloaded SiteSync and SlavePort	87
6.5.4	Architecture 4: multiple hardware partitions	88
6.6	Model outcome assessment	89
6.7	Summary	89
Chapter 7 Concluding remarks and future work		91
7.1	Achieved results	91
7.1.1	Development of a VDM-RT based methodology for Hardware/Software partitioning	92
7.1.2	Application in two case studies	92
7.1.3	Proposal of improvements to the VDM technology	92
7.2	Assessment of the VDM-RT methodology for Hardware/Software co-design	94
7.2.1	The notion of time	94
7.2.2	Time precision	94
7.2.3	The RTOS abstraction	94
7.2.4	Modelling communication	95
7.2.5	Code generation	95
7.2.6	Application in two case studies	96
7.2.7	Comparison with other methodologies	97
7.3	Future Work	97
7.3.1	Implementation of the AVB endpoint device proposed architectures	97
7.3.2	Guidelines for the transformation of VDM-RT to SystemC	98
7.3.3	Exploiting further the possibilities offered by Overture and VDM-Tools	98
7.3.4	Applying a heterogeneous modelling approach based on the VDM-RT language, 20-sim and SysML	98
7.4	Personal Learning Outcomes	99
7.5	Final Remarks	100
Appendices		111
A Terminology		113
B Servo case study VDM models		119
C AVB case study VDM models		177
D Representations and models used at different levels of abstraction		255

Table of Contents

E	MATLAB Based Methodology for Hardware/Software Validation	259
----------	--	------------

List of Figures

1.1	Thesis structure	8
2.1	Abstraction levels in a hardware modelling context.	11
2.2	Trade-offs in Hardware/Software co-design. Source: [Schaumont10]	17
2.3	Hardware/Software co-design process overview. Source: www.tu-dortmund.de	18
2.4	Level of details with regard to communication and time. Source: [Gajski&09]	20
2.5	Level of details with regard to communication and time. Source: [Black&04]	20
3.1	Venn diagram depicting SysML/UML relationship. Source: [SysML1.2]	25
3.2	Representation of POOSL classes under POOSL UML profile. Source: [POOSL]	29
3.3	Threading and step size in VDM. Source: [Lausdahl&10b]	32
3.4	Resource scheduling. Source: [Lausdahl&10b]	33
3.5	Xilinx Microblaze development tool-flow. Source: [Speedway11EDK]	34
3.6	Proposed approach by [Ou&05]	35
3.7	SATURN SysML based methodology. Source: [Mueller&10]	36
3.8	Proposed equivalences between SystemC and SysML. Source: [Mueller&10]	36
3.9	Overview of the SHE engineering design methodology. Source: [vanderPutten&07]	38
3.10	VDM based methodology. Source: [Larsen&10a]	39
4.1	Methodology overview.	42
4.2	Wrong usage of the hardware partition.	45
4.3	Correct usage of the hardware partition.	46
4.4	Application of a proxy class	47
5.1	Servo internals. Source: hitecrod.com	51
5.2	Typical servo control signals. Source: seattlerobotics.org	52
5.3	UML class diagram view of part of the sequential model.	54
5.4	Generated servo control signal.	55
5.5	Concurrent model simplified UML class diagram.	56
5.6	Distributed Real-Time model simplified UML class diagram.	58
5.7	Architecture of the software based solutions.	58
5.8	Deployment diagram of the software based solutions.	59
5.9	Generated control signal by the OS-less software only solution.	60
5.10	Generated signals by the RTOS based software only solution.	61
5.11	Generated signals by the RTOS based with priorities and the Hardware based solutions.	61
5.12	Architecture of the hardware based solution.	61
5.13	Deployment diagram of the hardware based solution.	62

List of Figures

5.14	IP block implementing the servo control functionality.	64
5.15	IP block deployed.	64
5.16	Software based solution profiling.	68
5.17	High pulse duration and correct signal period.	68
5.18	Generated control signals after 20 seconds of operation.	69
5.19	Detailed view of the generated PWM signal by the IP PWM block.	69
5.20	Overview of the generated PWM signal by the IP PWM block.	70
6.1	Protocol stack showing the involved protocols in AVB. Source: [IEEE-std1722-2011]	74
6.2	AVB and TCP/IP traffic managed by packet filtering. [Garner&07]	75
6.3	Overall structure of a Time Aware System. [IEEE-Std802.1As-2011]	76
6.4	A Time-Aware Bridge connected to a Grand Master. [IEEE-Std802.1As-2011]	77
6.5	Example network scenario.	78
6.6	Sequential model UML class diagram.	83
6.7	Context model UML object diagram.	84
6.8	Concurrent model UML class diagram.	84
6.9	Real-Time model UML class diagram.	85
6.10	Deployment diagram for the highly software based architecture.	86
6.11	Deployment diagram for the mostly software with offloaded SlavePort architecture.	87
6.12	Deployment diagram for offloaded SiteSync and SlavePort architecture.	88
6.13	Deployment diagram for the multiple hardware partition architecture.	89
7.1	Heterogeneous modelling approach.	99
D.1	Part of the memory map of an Audio Video Bridging (AVB) core interface. This memory map serves as an interface between the software and the hardware layers.	256
D.2	Used silicon area in a programmed FPGA. Layout information can be obtained through specific tools provided by the manufacturer. Source: xilinx.com . .	257
E.1	MATLAB Simulink used as a testing infrastructure. Source: [Simulink]	259
E.2	MATLAB Simulink combining an algorithm block and an HDL implementation. Source: [Simulink]	260
E.3	Combined test approach for HDL and Simulink simulations. Source: [Simulink]	260

Introduction

This chapter presents the work carried out in this thesis. We start with a description of the background of this work in section 1.1, the motivation that lead to its development in section 1.2, and the thesis goals in section 1.3. The methodology has been applied to two case studies, presented in section 1.4. A brief review of related work is presented in section 1.5. The typographical conventions used in this thesis are presented in the reading guide, section 1.6. Finally, the structure of the thesis is described in section 1.7.

1.1. Background

Today efficient embedded systems require hybrid solutions that are able to integrate hardware and software components. Complexity has grown exponentially over the last few decades, reaching sizes up to 1 million VHDL Lines Of Code (LOC) [Black&04], even in small sized Integrated Circuits (IC). On top of the hardware layer, one has to consider the software functionality implemented in the system. Software has challenges similar to the ones introduced by hardware but typically at a higher level of abstraction. Mastering both hardware and software is a must for all engineering teams working in the area of quality embedded systems.

An emerging challenge that results from the combination of hardware and software components in order to create systems, is the functionality allocation. This is also known as the *partitioning process*. The partitioning process is the stage in the design process in which it has to be decided whether a functionality has to be implemented in hardware or in software. This allocation/partitioning process is influenced by multiple factors like price, used silicon area, performance, meeting real-time deadlines and so on.

In order to master to incorporate such a variety of details in the design process, multiple authors agree on the possibility of using a heterogeneous modelling approach [Gajski&09, Waddington&06, Shaout&09]. The abstractions that are introduced in those models are as important as the modelling itself. Abstraction allows engineers to leave out selected details so as to attend to others, enabling a step-wise approach to system design. Some of the main results of using modelling techniques in the development process are a better level of problem understanding and a higher level of confidence in the developed solution.

Current methodologies in embedded systems design are derived from the System-Level-Design (SLD) approach. SLD pursues the paradigm of specify/explore/refine [Gajski&09]. It aims to start the development process by modelling at the highest level of abstraction and then perform refinement down to the final implementation. Hardware/Software Co-Design is a promising SLD

approach and a top-down methodology [Wolf03]. This approach is a model-driven engineering methodology, that aims to derive a system architecture from behavioural models. Deriving the system architecture implies performing the partitioning process that results in the definition of the hardware and the software blocks in a well-founded manner.

It would be fair to conclude that the application of modelling techniques play a major role in the development of new Hardware/Software systems. Models and abstraction are the only avenues to master the complexity present in today systems. Finally, traditionally it was preferable to implement functionality in software in order to lower the costs. Today's newer technologies and lower hardware prices are making possible to start considering hardware or software solutions as equal candidates. Approaches like Hardware/Software Co-design are gaining importance since they are focused on developing the functionality to be delivered, postponing the partitioning and functionality allocation process to a later stage in the process.

1.2. Motivation

This thesis work has been motivated by the author's interest in the Embedded Systems field. This interest was fostered in several courses during the Master's degree. The courses *Model driven development using VDM++ and UML* and *Hardware/Software Co-design* deserves special mention. The Model driven development course presented the way concepts like abstraction and formal modelling can be applied in the development of systems. Through the application of abstraction a complex problem can be boiled down to its essence, being easier to be approached and solved. A precise and abstract representation of a system can be created by the use of formal models. This representation allows the analysis and lays a solid basis for the implementation stage. Hardware-/Software co-design makes very extensive use of models in order to represent and analyse different architectural solutions. The applied models in this field can be at different abstraction levels, and can range from general system models for initial system specification to concrete subcomponent representations.

There are already several modelling technologies for Hardware/Software co-design in the market. Every modelling technology presents an associated methodology that describes its application throughout the development process. Some of these technologies are too generic and, even though they can be fitted for the purpose of communication and overall analysis, a thorough study of the problem becomes difficult to perform (e.g. SysML, UML profiles). On the other hand, there are very specialized technologies that tackle concrete problems in a proficient way, but are weak on representing more generic scenarios (e.g. Matlab Simulink). It is the author's belief that there is a need for a modelling workflow that takes the most beneficial aspects of both approaches and provides a way to conduct hardware/software design space exploration.

1.3. Thesis goal, approach and scope

This section describes the Master's thesis goal, the approach used and the scope of this work.

1.3.1 Goals

The main goals of the thesis are:

Approach

1. **To propose a general purpose methodology based in the VDM-RT modelling language in order to support the design decisions in the Hardware/Software co-design field.**
2. **To propose improvements and extensions to the VDM-RT language and the modelling tool Overture, seeking to contribute on the creation of a modelling tool that can be applied in a more specialized field.**

An additional personal learning goal of this thesis, as important as the previous ones, is to improve the author's skills and knowledge in the VDM-RT modelling technology and in the Hardware-/Software Co-design field.

1.3.2 Approach

The methodology applied in this thesis is structured in two phases, each containing several points of action.

Phase 1: Study, analysis and reflections

Reading and research activities that will allow the author to: evaluate the need for new methodologies and draw the requirements for new methodologies in the case they are needed. The result of this phase will allow to build the rationale behind the work that will be carried out during the rest of the thesis. This phase is decomposed in the following steps:

- Research on the state of the art of modelling techniques applied in system-level design.
- Exploration of current processes and workflows used in hardware/software co-design. This research includes both academic readings and interviews with experts in the field.
- Research on the current challenges in the hardware/software co-design field.
- Additional study on the FPGA technology.

Phase 2: Elaboration

The elaboration phase covers some of the methodological needs detected in Phase 1 by the design of a VDM-RT methodology for Hardware/Software co-design. This phase is decomposed in the following steps:

- Development of a VDM-RT based methodology in order to provide partial support in hardware/software co-design decisions.
- Application of the proposed methodology in case studies in order to validate its performance.
- Pragmatic analysis of results and evaluation of the proposed VDM-RT based methodology.

1.3.3 Scope

This work will be focused on:

- Proposing a methodology that supports the process of making partitioning decisions during the development of hardware/software systems.
- Focusing on partitioning decisions that can be reached through the analysing of real-time deadline accomplishments, interaction between components, and approximate system behaviour modelling.

This thesis does not have as an ultimate goal the implementation of industrial solutions for the problems presented in the case studies. Actual incorporation of the proposed improvements to the tool Overture or the modelling language VDM in any of its flavours are out of the scope of this thesis work.

1.4. Case studies

The proposed methodology has been applied in two case studies. A core part of the thesis has been the evaluation of different partitionings with regard to real-time deadlines. The first case is conceptually simple with a single real-time deadline, aiming to show how the methodology should be applied. The second case is based on a real industrial application and with multiple real-time constraints and several possible partitionings. A brief description of these cases is provided below, for further information refer to chapters 5 and 6.

- **Development of a servo controller:** The servo controller case will present how different implementation strategies can be modelled by the use of VDM-RT. This case shows the application of the methodology in a single-process scenario with real-time deadlines. This is an introductory case study based on the workshop [Speedway11SDK].
- **Development of an AVB endpoint device:** Audio Video Bridging is a technology that aims to offer real-time multimedia content streaming over Ethernet network. This technology requires the implementation of protocols for time synchronization (802.1AS), traffic shaping (802.1Qat) and stream reservation (802.1Qav). The case is ideal for testing the performance of the methodology in a time-critical, multi-process setting. The modelling focus will be set on the time critical aspects of AVB. This case study has been provided by Bang & Olufsen.

1.5. Related work

Before starting considering a new approach for Hardware/Software co-design, a review of the current modelling techniques, methodologies and challenges has been performed. In this section, only an overview of the considered work is presented. Additional details in existing approaches will be presented in chapters 2 and 3.

Formal modelling processes for Hardware/Software engineering

[Huang&04] proposed the **Software Hardware Engineering** methodology, based on the POOSL modelling language. This methodology is based on a formal object-oriented modelling language. It has been applied in a number of industrial projects and several tools have been developed even allowing model compilation. Unfortunately this area of research seems to be inactive. In his thesis, Verhoef introduces a real-time extension for the formal modelling language VDM [Verhoef05]. This extension provides the necessary constructs for the evaluation of different system architectures in a distributed real-time deployment. Additional research and development effort has been spent on this and a concrete implementation of this is in the modelling tool Overture. A complete development process entirely based on the VDM method is proposed in [Lausdahl&10b, Larsen&09]. This methodology makes use of the VDM-Specification Language, the VDM-PP Object Oriented modelling language and, finally the VDM-Real Time extension. This process makes use of formal model during the whole project life-cycle. The tool support

for this process is provided by Overture [Larsen&10b] and by VDMTools [VDMTools]. There is no previous work related to the application of the VDM methodology in the Hardware/Software Co-design field. The current research in VDM is working toward the integration with other modelling platforms in fields like mechanical co-simulation [DESTTECS09] or System-of-Systems [COMPASS11].

Specialized processes for FPGA based systems

Xilinx presents a development process for FPGA-based system design in [XilinxToolFlow11, Speedway11EDK]. This workflow is highly tool-oriented and slightly covers the problems associated to hardware/software co-design of embedded systems. Simulink provides a sophisticated modelling framework that can be used in many modelling activities [Simulink]. One of these activities is the modelling of FPGA-based systems. Ou et al. [Ou&05] makes use of Matlab/Simulink in order to co-simulate different hardware partitions. The same authors propose an interesting contribution where Matlab models are applied for energy estimation consumption in FPGA soft-processors [Ou&04]. The existing literature is focused on using Matlab as a specialized tool in order to create models that will help on taking concrete partitioning decisions. Little effort has been spent on relating these modelling activities with a complete methodology to be applied for complete system development.

Combined SystemC - UML profiles based processes

Mischkalla et al. and Prevostini et al. propose the combination of SystemC with UML profiles in order to create a sustainable modelling-implementation process through the development process [Mischkalla&10, Prevostini&07]. Such an approach is interesting since two modelling technologies are combined. This heterogeneous application is taking the most beneficial aspects of both technologies and mitigates possible shortcomings. However, UML profiles can be limited when it comes to model execution and SystemC models might be biasing partitioning decisions towards software (SystemC is a C++ class library). [Mueller&10] describes the approach followed in the ongoing EU FP-7 project SATURN (SysML Based Modelling Architecture exploration, simulation and synthesis for complex embedded systems). In this approach an addition to Artisan Studio is used to create SysML models that can be synthesised. The generated code can be co-simulated in a virtual platform for system verification. Finally, generated code can be deployed in the actual platform.

The proposed methodology for Hardware/Software Co-design tries to approach the hardware partitioning problem, that is not explicitly addressed in the existing VDM-RT development process [Larsen&09, Lausdahl&10b]. The proposed methodology benefits from several ideas proposed in the previously cited papers.

One of the most relevant features of the VDM-RT based methodology for hardware/software co-design is that it is not specialized to address particular problems. This is beneficial from the system perspective, as it is better suited than the specialized processes for FPGA-based system. However, when it comes to solve particular implementation details [XilinxToolFlow11] and [Simulink] might be more adequate.

Finally, the combination of graphical notations applied in the combined SystemC - UML profile based processes are interesting and the synthesis possibilities outstanding. Note that this has been reached after years of development and considerable investments, therefore it will not be fair to compare the methodologies considering the state of the tools. Considering only the modelling methods used, it could be concluded that SysML might lead to an informal use of the modelling technology, implying imprecisions and modelling errors. SystemC is too close to the implementation language C++ and might alter the focus of the modelling and bias the system engineer to

more software-based implementations. On the other hand VDM-RT is far from final implementations of the system, since it is an abstract modelling language by construction. A more in-depth comparison of these approaches will be provided in chapter 3.

1.6. Reading guide

This thesis makes use of the following conventions:

References:

Articles, books, technical reports, figures and anything that is not the author original work is referenced. The references are placed between square brackets containing the surname of the primary author and the year of publication, e.g. [Brooks86]. In the case there are several authors for the same resource the character & will be placed between the surname of the primary author and the year of publication, e.g. [Black&04]. Finally, in the case the author has published several articles during the same year, a letter will be added following the year of publication, e.g. [Lausdahl&10b].

Emphasis:

Concepts, words or sentences with special relevance are emphasized in the text by the use of *italic*.

Guidelines:

General instructions and modelling advice are given in the form of guidelines, that present the following style:

Guideline 9: Each hardware candidate should be deployed on an individual CPU configured as a hardware block.

Keywords and language instructions:

C, VHDL and VDM keywords are shown in **typewriter boldface**. Other instructions are shown in plain typewriter font.

VDM model listings:

VDM models are placed in listings that present the following style:

```
1 private generateCycle: () ==> ()
2 generateCycle() == ();
```

VHDL code listings:

VHDL source code snippets are placed in listings that present the following style:

```
1 control_counter_process : process (clk, rst)
2 begin
3 end
```

C code listings:

C source code snippets are placed in listings that present the following style:

```

1 void rc_servo_set_servo_position(int baseaddr,
2                                 int servo_number,
3                                 int position);

```

1.7. Structure

This thesis is organized in seven chapters, of which this introduction is the first. Figure 1.1 presents the thesis structure in a graphical manner.

Chapters 2 and 3 summarize the theoretical foundation of this work. These chapters will provide the necessary concepts in System Level Design, Hardware/Software Co-design, Workflows and Modelling.

- Chapter 2 introduces the concept of abstraction and its application in the Hardware/Software engineering field. Special emphasis is made in the importance of the partitioning decisions during the analysis and design phases of the project.
- Chapter 3 gives an overview of the most relevant modelling languages used in system level design. Together with the languages, workflows that make use of them in order to model and study Hardware/Software systems are reviewed.

The following chapters present and apply in practical cases a new methodology for design space exploration in Hardware/Software co-design:

- Chapter 4 presents the VDM-RT methodology proposed by the author. Concrete guidelines and modelling advice will be given so it can be applied by a modeller with minimal knowledge of VDM-RT. A critical view on the methodology will be presented.
- Chapter 5 introduces the case study on servo control. The proposed methodology in this thesis will be applied in a case that is easy to understand by a reader with minimal knowledge on hardware design. The methodology will be applied from the analysis and design stage down to the final implementation of two different architectures. This case study will give the possibility of validating the hypothesis derived from the VDM models against real implementation details.
- Chapter 6 introduces the second case study of this thesis which is based on the development of an AVB endpoint device. The system under consideration in this case study is complex and involves several protocols. The VDM-RT methodology will help to detect critical aspects and to provide information in order to make design decisions. An actual implementation in this second case study has not been performed due to time limitations. This case study has been proposed by the commercial electronics manufacturer Bang & Olufsen.

This thesis is concluded in chapter 7, which summarizes and analyses the achieved results. It also establishes areas in which further work could be performed.

This thesis contains five appendixes. Appendix A provides a list of definitions and abbreviations of hardware/software terms used in this thesis. Appendixes B and C present the produced

Chapter 1. Introduction

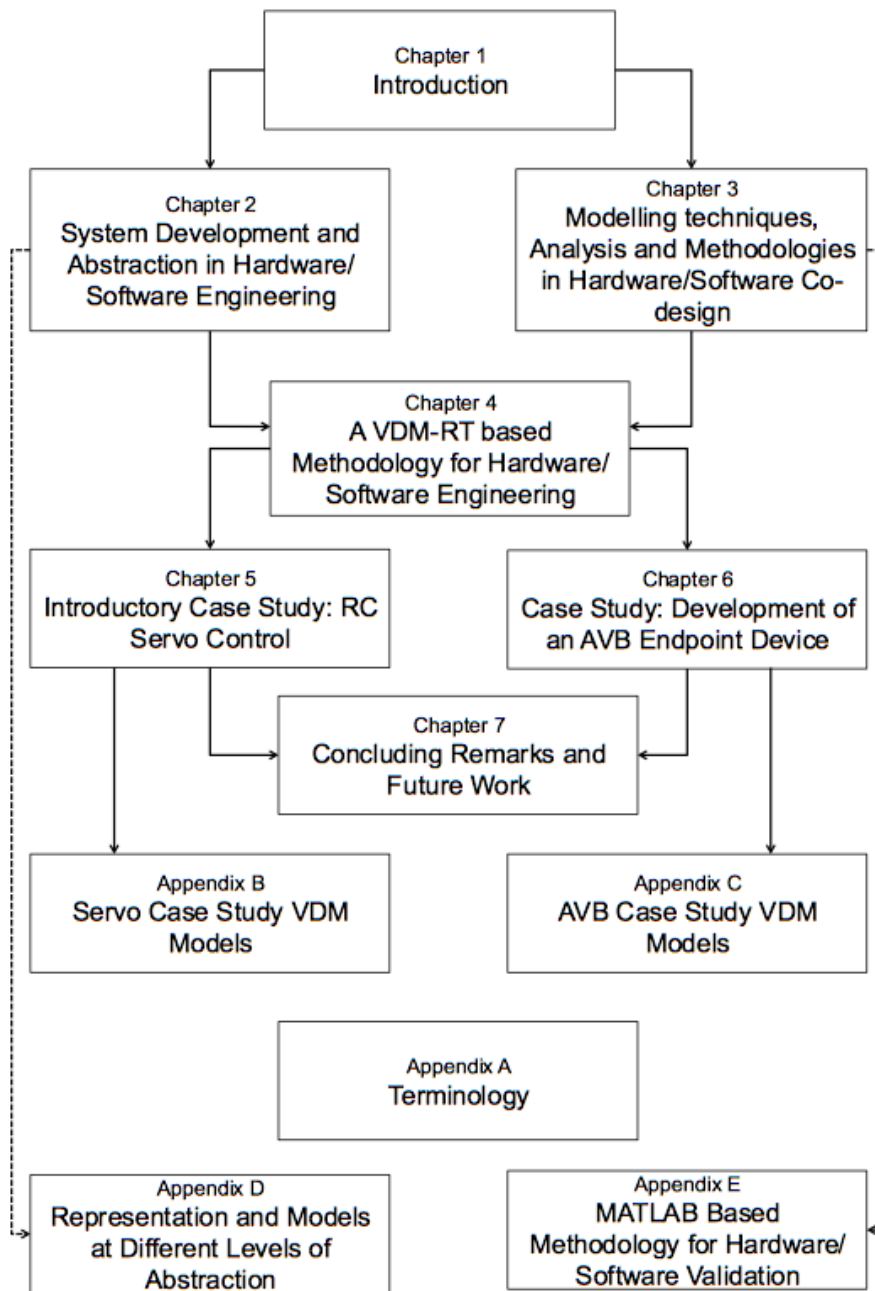


Figure 1.1: Thesis structure

VDM models for the servo and the AVB case studies respectively. Additional details on representations and models used at different levels of abstraction are given in appendix D. Finally, more information about the MATLAB based validation methodology is provided in appendix E.

System development and abstraction in hardware/software engineering

This chapter introduces system development and complexity, presenting the rationale for new methods and tools in section 2.1. Abstraction levels that can be used in a Hardware/Software Engineering context are described in section 2.2. The approach System-Level Design and the modelling languages used under it are described in sections 2.3 and 2.4. The concrete System Level Design approach hardware/software co-design is presented in section 2.5. Finally, the chapter is summarized in section 2.6.

2.1. Introduction

Before presenting methods, tools and system related abstractions, the term system itself requires to be defined. According to the Systems Engineering Handbook [INCOSEseh10] a system is defined as:

"An integrated set of elements, subsystems, or assemblies that accomplish a defined objectives. These elements include products (hardware, software, firmware), processes, people, information, techniques, facilities, services and other support elements.

Besides the system engineering perspective present in the development, that is considering even facilities and services, the element under study in this thesis is what is defined as *product*. Both, [NASAseh10] and [INCOSEseh10] qualify *hardware*, *firmware* and *software* as crucial elements to produce system-level results.

Systems are developed by carrying out different processes which, according to [INCOSEseh10], can be classified as organizational, technical and managerial. The development processes proposed in this thesis work falls under the category of *technical processes*.

The purpose of the lines above is to contextualize the methodologies applied within the technical processes, that will be discussed in chapters 2, 3 and 4. The ultimate aim of these processes is to support and perform the development of a product, composed by *hardware*, *software* and *firmware*.

Hardware/Software systems are composed by a wide variety of elements and components of different nature, presenting a growing complexity. This complexity has increased exponentially over

the last decades and traditional approaches, based on simple methodologies and the engineer's experience, are not always enough to tackle the project complexity. [Black&04]. This kind of situation leads to cause *analysis-paralysis*¹ in the system development process or *poor decision making* overlooking at critical details, something that may flyback later on in the project life-cycle. With the purpose of dealing with this complexity, [Patel&04] proposes to rise the level of abstraction in system design and to use different methods and tools in system development. A similar argument is hold by [Black&04]. Gajski in [Gajski&09] presents that it is needed a systematic and well-founded approach to design combined hardware/software solutions, that supports the specification phase and provides solid arguments to perform design decisions such as system hardware partitioning.

The current approach in System Design Methodologies is, according to [Gajski&09], and from the early 2000s, *Specify, Explore-and-Refine methodology*, which introduces the so called *System-Level (SL) perspective*. This approach is primarily using the concept of abstraction in order to define different levels. Abstraction levels are enabling an stepwise system development, which helps on facing a challenging project, in which the development team has limited experience or is specially complex.

2.2. Abstraction levels overview in a single system

Before going in detail on abstraction levels, the concept of abstraction itself deserve special consideration. There are different definitions explaining what abstraction is focusing on different aspects. Kramer in [Kramer07] is presenting abstraction as the process of *removing detail to simplify* and as a *generalization to identify* the common core. The first definition, which is the most connected with this thesis work is a quote from Kramer:

"the act or process of leaving out of consideration certain properties of a complex object so as to attend to others."

In the system development context, abstraction is playing a key role allowing the engineers to tackle the complexity inherent in hardware/software systems. As Kramer explained and applied in this context, this tool will allow the designers to consider some of the properties that should be present in a system without considering others, that will be taken in account further in the development process. Therefore, abstraction is a enabling tool to describe the system at different levels and to perform different kind of analysis.

The way from a system described in terms of inputs and outputs (blackbox approach), to a system described with a fully detailed hardware implementation is organized in several layers of abstraction. All of them are relevant during the development process, providing details in different implementation areas. Abstraction is therefore, mainly, a tool to deal with the complexity and the amount of details present in a system.

In a combined hardware/software context, multiple authors ([Gajski&09, Black&04, Sudhakar05]) agree on six abstraction levels for single-device based systems. The levels that this thesis work is covering are:

¹Common term to refer to over-analysing a situation, causing no productive results and no progress in decision-making.

Abstraction levels overview in a single system

Algorithmic Level: The algorithmic level is focused on the behaviour the system has to implement. At this level, no description of the hardware implementation is made. The effort is focused on describing the information the system is receiving, how this information is processed and which outputs should be provided. At this level it is possible to analyse the software logic. Relevant aspects could be concurrency, protocols, policies or data consistency.

Interface Level: The interface level incorporates details about the interface between the functionality implemented in hardware and the software components. This level is commonly referred in the literature as the "programmers view" At this level the hardware partition is clearly defined. Some of the analysed aspects are data availability, data format and protocols.

Transaction Level: The transaction level view adds timing details to the interface. It offers time related details that can be useful to estimate real-time deadline accomplishment. Appropriate real-time performance at this level is necessary condition but not sufficient. It is also known as the "timed programmers view", and models at this level are named Transaction Level Models (TLM). At the transaction level approximate time details are used to analyse timing details related to the hardware/software interface.

At a lower level of abstraction we can find the *Cycle Accurate Level*, the *Register Transfer Level* and the *Silicon Level*. These levels are not relevant for this thesis work and further information can be found in [Gajski&09].

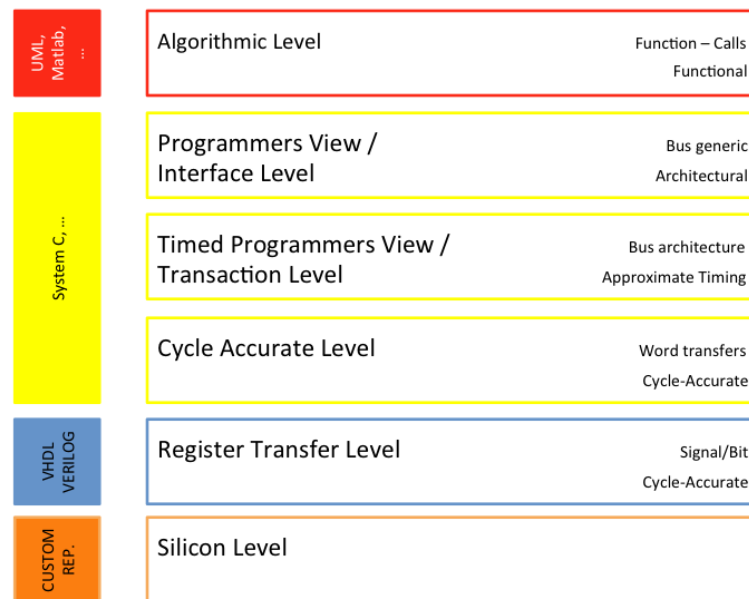


Figure 2.1: Abstraction levels in a hardware modelling context.

Figure 2.1 is detailing the abstraction levels discussed above. Note that the targeted levels by this work are at the top of the abstraction stack. The algorithmic level is often stratified in several sub-layers and varies in complexity. This last issue has to do more with the embedded software development field rather than with the partitioning process, therefore this thesis will not elaborate further in that area.

More information about the different representation and models that can be used at each abstraction level can be found in appendix D.

2.3. The System-Level perspective

The system-level design approach is aiming to start the design process at the highest abstraction level of the ones explained above. From this level, the system is represented by an executable specification that represents the behaviour (*Specify*). Additional models are incorporated to this holistic view in order to illustrate different facets of the system (*Explore*). These models can provide details on design decisions and serve as a tool to evaluate different configurations. Once exploration by the means of these models is concluded, additional implementation details are added (*Refine*). Each model refinement will be carried out after taking more design decisions, therefore incorporating more details about the final system. A relevant characteristic of SL design is to reuse components and Intellectual Property (IP) blocks whenever is possible, minimizing the development effort.

The system level design approach follows a sound engineering procedure applied during many years in multiple fields. SL starts with a initial specification, followed by possibilities exploration and finally model refinement down to actual implementation [Gajski&09].

2.3.1 Different approaches in system-level design

There are three main system-level design methodologies that make use of the approach Specify-Explore-Refine introduced above [Shaout&09, Gajski&09]:

Hardware/Software Co-design: also known as top-down methodology or Model-Based Design.

This methodology pursues to model the behaviour of a certain system in order to derive an architecture from it. Deriving this architecture implies cooperative design of hardware and software.

Platform-Based Design: this approach allocates the functionality implementing the behaviour to a previously defined architecture.

Component-Based Design: also known as bottom-up methodology. This methodology starts by creating the building blocks at the lowest level of abstraction. These blocks will be used in the creation of the immediate superior layer in the implementation. All blocks are organized at libraries present at the different abstraction levels.

2.3.2 Application of system-level design approaches

The application of System-Level design methodologies depends on the encountered situation. *Platform-Based Design* is the required approach in the case a certain platform has to be used. For instance, in the case an update over a certain system has to be performed and the architecture for this system has been already defined. The system might be featuring several processing units, so the issue considered in a platform based design approach is in which out of these processing units the functionality has to be allocated.

Hardware/Software co-design will support the development of an optimized platform, supporting this task with metrics on different factors. The use of this flavour of System-Level design approach is ideal for development of new products, in which a higher degree of flexibility is present. Design space exploration is fostered in this approach.

Finally, *Component-Based Design* requires an extensive creation of components library at each abstraction level. These components have to be created, ideally, before moving to the superior layer. It is an arduous task, if not impossible, to make a prediction of exactly which components

are to be used in a certain product, therefore there have to be components implementing many different functionalities which are not going to be used in a certain system. As an advantage, the bottom-up methodology facilitates the creation of re-usable building blocks, clearly allocated in an abstraction level. This approach could be adequate for a large hardware manufacturer, which will possibly reuse components in later developments.

From now on, the focus of this work will be the system-level design approach pursuing Hardware/Software co-design.

2.4. Modelling languages in system-level design

The essence of System Level (SL) design, as expressed by [Shaout&09] could be resumed as:

"System Level Design maps a high-level abstract specification model of an entire system onto a target architecture."

Besides Shaout et al., many authors agree that speaking about System-Level Design without considering modelling would be impossible. Since the moment in which the concept of abstraction is present in this way of designing systems, it is clear that certain kind of system representation is going to be present. This directly implies that models are going to be used, which lead to consider two issues: a) The kind of languages needed to make such a representation and b) The level of and kind of details that should be present at this initial, abstract representation. This issue has been widely discussed in the literature.

At this point, more relevant than reviewing concrete languages, it is more interesting to consider which possibilities a modelling language should offer in order to be useful in a System-Level Design context. According to [Shaout&09, Gajski&00, Niemann98] a modelling technology used to specify a system should incorporate facilities to represent the following elements/properties:

Hierarchy: as explained above, increasing size of systems makes necessary to introduce structures in order to organize system behaviour and architecture. These structures typically have subordinated components solving part of the problem that they are responsible for. Making a proper representation of a system structure is a basic feature.

State transition: system operation can always be described as a state machine. Depending on external inputs (current, previous or both), the state can be in different operational modes or states of execution. Some embedded systems are purely implemented as state machines, its use is a common practice at both Software and Hardware level.

Concurrency: in many cases a system is executing several processes or threads at the same time. There is a leap in complexity from single-task based solutions to multi-task ones, both at the implementation and specially at the debug stage. The modelling language should be a tool to support this situation and foresee possible conflicts and concurrency problems.

Synchronization: The need for synchronization modelling support is derived from concurrency. Race conditions, deadlocks and other concurrency related problems can cause complete system malfunctioning. Like concurrency problems, these cases can be extremely difficult to detect. In some situations, problems are detected when the system has been already deployed. With the introduction of multi-core processors, synchronization must be guaranteed even at the hardware level. A modelling language supporting concurrency must support synchronization in order to be useful in the multi-task domain.

Interrupt handling: interrupts are common mechanisms in embedded-systems, often used to signal events. These kind of events may require the interruption of the current task and, therefore, transition to an alternate state to process them. Some embedded systems are purely interrupt driven. An modelling language used to model embedded systems must be able to represent this kind of situations.

Timing: many embedded systems present soft/hard deadlines in its operation. Not satisfying these deadlines could imply complete system malfunctioning. Modelling constraints of this kind, specially in concurrent systems is a desirable property.

Constructs to support formal verification: it should be possible to express certain system properties so logic and discrete mathematics can be used to prove that they are present in the system. While powerful, formal verification is sometimes far from the embedded engineer mindset. Strong tool support is required to make formal verification usable in an industrial setting.

Dynamic behaviour: in the case the system is making use of dynamic scheduling in order to coordinate software operation, the modelling language, ideally, should provide a way to represent and analyze it. Dynamic scheduling is often implemented by the use of a Real Time or ordinary² operating system.

In [Huang&04], the authors rise the level of abstraction and require different features that a modelling language should present:

Adequate expressive power: the modelling language has to be capable of expressing complex properties, behaviour and structure.

Platform-independent semantics: the semantics of the language shall not be bounded by a certain technology. The use of the language should not be restricted to a concrete platform.

Operational semantics: this property refers to the *executability*. By executing a model, simulation and/or model-checking techniques can be applied in order to detect inconsistencies and prove correctness.

Modularity support: as explained above, the possibility of expressing the internals of a component and its interfaces is a must.

Automatic and correctness-preserving transformation: the system-level design model should be a *golden model*³, that will be refined until the actual system implementation is achieved. During this process, a large number of refinements and transformations will be applied to the initial model. By using a tool that is automating these transformations and, furthermore, ensuring that certain properties are satisfied, the development of the system is easier.

It is difficult to find a modelling language that incorporates all the features suggested by Huang, Gajski and Niemann. A compromise solution must be reached when deciding on a concrete language. Further details on concrete languages and application areas will be introduced in chapter 4. This requirements are giving an idea about the kind of details that should be present at a System-Level Design model, or at least considered in an initial specification.

²Non-Real Time Operating System

³Common term to refer to a model that is used as a reference during the whole development process.

2.5. Hardware/Software co-design

Nowadays embedded systems are often implemented in heterogeneous architectures, that combine dedicated hardware blocks, processors and software components deployed in processors⁴.

The possibility of using different hardware and software blocks opens a wide range of possibilities when it comes to the architectural design stage. The challenge in this situation is to decide whether a certain behaviour should be implemented as a software or as a hardware component. The Hardware/Software co-design discipline provides the tools and the methods to generate architecture from behaviour by adding implementation details to the design.

Before digging in to some of the most outstanding aspects about this discipline, it is worthwhile considering definitions from the literature:

[Schaumont10] defines this discipline as:

"Hardware/Software co-design is the design of cooperating hardware components and software components in *a single design effort*."

The same author remarks as well:

"Hardware/Software co-design *focus on the partitioning and design* of an application in terms of fixed and flexible components."

[Wolf03] remarks:

"... hardware/software co-design tries to increase the predictability of embedded system design by providing analysis methods that tell designers if a system *meets its performance, power and size goals* and synthesis methods that let researchers and designers rapidly *evaluate many potential design methodologies*."

Finally, [Shaout&09] characterize Hardware/Software co-design as follows:

"... also referred to system synthesis, is a top-down approach. Starts with system behaviour, and *generates the architecture from the behaviour*. It is performed by gradually adding implementation details to the design."

These four quotes jointly give a good overview of what Hardware/Software co-design is: "a single design effort", what the main focus is: "the partitioning and design", which criterion are relevant in partitioning: "performance, power and size goals" and, finally, how the system is produced: "architecture derived from behaviour". The following sections elaborate on the details outlined by these three authors and presents FPGAs as platforms for using this methodology.

2.5.1 Trade-offs and impact of Hardware/Software co-design

Every system specification presents a certain collection of targets that shall be fulfilled in order to produce a successful implementation. Some targets, as they were mentioned by [Wolf03], are specified in the form of requirements that *must* be satisfied. Other targets, are design variables that can be optimized in order to produce a higher quality system. In order to consider a design target as fulfilled, and in order to optimize its variables, it is needed to apply a specific criteria. By

⁴In this case "processor" is a common term to denote to a "microcontroller", a "Digital Signal Processor" or a "microprocessor".

applying it, it is possible to compare solutions and determine which one is valid/better. Connected criterion with opposite gradients in its optimization function are known as trade-offs.

Considering a trade-off situation; a change on one of the involved variables might have an impact on the rest. For example, raising the frequency of the processor speed up the computation therefore, the performance. The higher the frequency the higher the energy consumption. In some cases the higher the temperature of the silicon.

The challenge here is to compromise in order to find the optimum balance in these trade-offs. The most relevant criterion, that might imply trade-offs from different perspectives in the design space, are explained below. These criterion is proposed by the author taking inspiration from [Chen&09] and [Schaumont10].

From a technical point of view, the most important variables that might present trade-offs are:

Performance: The level of abstraction in which a certain solution is implemented has a direct connection with its performance. By construction, the hardware layer is at a lower level of abstraction than the software layer. The existing overhead in a hardware functionality implementation is practically null. A hardware implemented functionality presents a higher degree of performance in *all* the cases. A hardware implementation might be more expensive to create in terms of time and material costs.

Flexibility: Functionality implemented as a software component is more flexible in terms of re-usability and modification. A hardware block is typically suited for a single purpose and can hardly be reused for a different purpose.

Energy efficiency: energy constraints are present nowadays in many systems. Some systems require to be energy efficient so their operational life is longer e.g. a battery powered multimedia system. Others require to be energy efficient in order to be environmentally friendly e.g. a class-A home appliance. Finally, due to pure physical reasons, the more energy an electronic device is consuming the more heat is going to be dissipating. This might lead to performance decrease or even compromise system availability. A system that is operating at a higher frequency might perform better, but it will consume more energy.

Complexity: The heterogeneity present in a System-On-Chip (SoC), presents a challenge itself. Considering a hardware block in isolation, several thousands lines of VHDL code have to be used in order to implement a functionality that might require hundreds in the C programming language. This means that, a hardware implementation is harder to design, create and maintain. Furthermore, debugging facilities in software are more simple than the ones available in hardware. The latter kind often implies a combination of external equipment, in-system logic analyzers⁵, ... On the other hand a hardware solution generally presents better performance than a software solution.

Co-design decisions have an impact as well from a managerial point of view. A certain architecture will present particular time and cost demands:

Engineering hours cost: Development effort is one of the most relevant issues when evaluating different design possibilities. The amount of man-hours that will be spent in the development depends on the available time-budget assigned to the project. A heavily based hardware solution will have a considerably higher demands on engineering development hours than a combined hardware/software approach.

⁵An example of in-system logic analyzer is Xilinx Chipscope. This tool enables the developer to analyse the logical signals inside the FPGA. For additional details refer to www.xilinx.com.

Material cost: Material cost is one of the most relevant recurring design variables. Even minimum savings per unit lead to an important revenue in the long run. This might be in conflict with the performance or the energy efficiency variables.

Development time availability: Current product developments are characterized by the short time-to-market. Reusability of hardware blocks whenever is possible, integration of IP cores or partitioning are some of the factors that will make certain architecture feasible within the available time frame for the development.




Figure 2.2: Trade-offs in Hardware/Software co-design. Source: [Schaumont10]

The technical criterion could be refined even further and consider more particular details. These details depend on the problem under study. In [Brogioli&06], the partitioning is performed considering mainly the real-time deadlines for a 3.5G mobile receiver. In this case the data and computational related aspects are especially relevant. The goal in this case is to derive an architecture that fulfils the real-time deadlines. The approach is to improve the performance of the system. The performance sub-criterion used in this partitioning problem were:

Data spatial location: Data can be stored in a number of different hardware elements in a certain architecture. Latency, access and setup times or bus usage make a difference when it comes to read or write data. Where to place the data, besides the used memory technology, will affect the time performance of a concrete system architecture.

Data level parallelism: Parallel data availability can be exploited from an implementation point of view. FPGAs and DSPs to some extent are strong on receiving a parallel stream of data reducing data input latency in the system.

Task level parallelism: An implementation in which multiple elements are solving a complex problem in a parallel way is considerably more efficient than a serial processing approach. When combined with a high degree of *Data level parallelism*, a parallel implementation might present an outstanding level of performance. However, not all computational problems are suitable to be solved by parallel computing. For example the image processing field present numerous algorithms suitable for parallel computing. On the other hand, the algorithms used in a network interface are not the ideal case for parallel processing.

Computational complexity: An algorithm running in a DSP or CPU might present bottlenecks that can be detrimental to the overall performance. The solution in many cases comes in the form of *hardware off-loading*⁶, which implies moving the worst performing algorithms executing in the host processor to a hardware implementation in a FPGA. The criteria to determine which algorithm is performing best is not always the same and, in many cases, it is ad-hoc. In some cases the validation criteria can be time deadlines accomplishment, in other cases throughput, etc...

⁶A concrete hardware off-loading case will be presented in chapter 5

These aspects might not be critical or not even present in all applications. The conclusion that can be drawn is that, besides the general issues described above, particular partitioning sub-goals might arise in a certain project. This presents, a challenge for a co-design methodology. The applied techniques should be flexible enough in order to incorporate the *heterogeneous details* and *particular* considerations that each problem might present.

2.5.2 Overview of the Hardware/Software co-design process

There are several methodologies and techniques that can be applied in order to do Hardware/Software co-design. All of them are pursuing the objectives described above. Even though they might differ in techniques, artefacts or performed analysis, all of them follow the overall structure presented in figure 2.3.

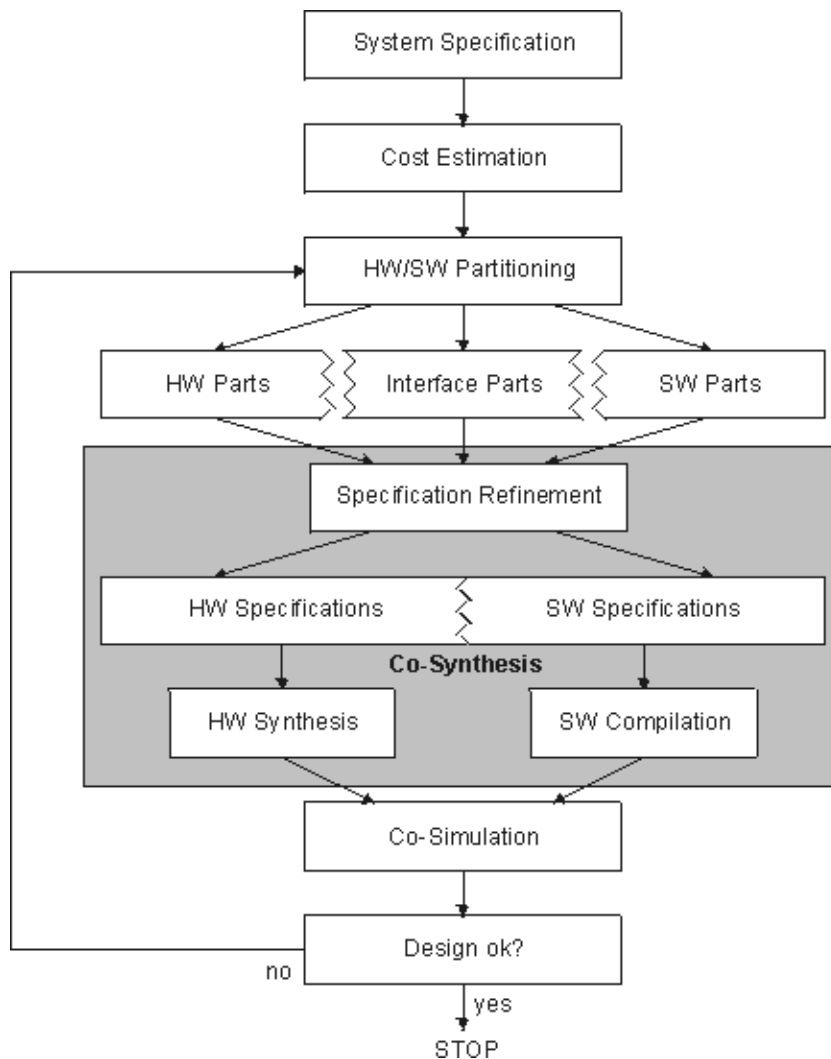


Figure 2.3: Hardware/Software co-design process overview. Source: www.tu-dortmund.de

The system design stage starts with the phase *system specification* in which functional and non-functional requirements are defined. In *cost estimation*, a cost factor for each design alternative is determined. This cost factor enables the system architects to evaluate the design alternatives against a set of criterion previously defined. These criterion might be similar to the ones introduced

above. There are different models that can be used to perform cost estimation. Ad-hoc models, system models or previous experience are the most typical ways of carrying out this phase. This phase is where the design space exploration takes place. Once the optimal solution given the situation has been determined, the *Hw/Sw Partitioning* takes place. This phase will determine the system architecture, and will result on the definition of *Hardware parts*, *Interface parts* and *Software parts*. The next stage in the design is the co-synthesis, in which the Hardware and Software blocks are created in parallel. A key role in this stage is presented by the *refinement towards hardware specifications and software specifications*. The result of the co-synthesis is a model that can be co-simulated, so hardware and software blocks can be run together in a software platform. Results can be validated and verified against the initial requirements. In the case the validation is successful a bitstream for a concrete FPGA and a binary for a certain processor can be generated and deployed on target. Further testing activities on a real⁷ hardware platform are performed and finally on the actual platform. The *Hardware/Software co-synthesis* and the *Real and Actual hardware testing* stages are out of the scope of this thesis.

2.5.3 The role of modelling in Hardware/Software co-design

As explained in sections 2.3 and 2.4, modelling techniques are fundamental tools in the System Level Design approach. Since Hardware/Software co-design is a model-driven engineering approach, the relevance of modelling is even greater. Many authors highlight the importance of models of computation and its application through the design process [Gajski&09, Black&04, Schaumont10, Shaout&09, Chen&09].

Two concepts are crucial in order to understand models in System Level design and, in particular, in Hardware/Software Co-design: a) Abstraction and b) Refinement. Both of them have been already explained in section 2.1. When considering models in this context, and according to Gajski et al. and Black et al. in [Gajski&09, Black&04], it is important to keep in mind that:

- The less abstract a model is, the closer to the final implementation it is;
- refinement is performed when lower level details are incorporated to a higher abstraction level model;
- the added details are relevant from the communication or from the computation (functionality) point of view and
- the precision of the notion of time can be improved by refinement.

The initial modelling activities start with the creation of a Specification Model (A), this model is named by [Black&04] as System Architectural Mode (SAM). A SAM model does not incorporate any notion of time. The main outcome is the *causal ordering between processes*. The causal ordering allows to describe what happens before and after considering certain event, in other words ordering of events and actions. The next refinement considering computation is the *Timed Functional Model* (B). At this step time annotations of certain computations and estimated delay results are incorporated. The first model incorporating an approximated notion of time in both communication and computation aspects is a Transaction Level Model (TLM) (C). From this point it is possible to refine the models in terms of communication or in terms of computation. A TLM model that is reaching a Cycle Accurate representation is called Bus-Cycle-Accurate

⁷The term real hardware platform refers to an FPGA development board. The term actual hardware platform refers to the final hardware device, the electronic platform that will be mounted inside the product to be developed.

Model (BCAM) (D). On the other hand, a TLM model refined in terms of computation is called Computation Cycle Accurate Model (CCAM) (E). When both BCAM and CCAM models are combined a Cycle-Accurate-Model (CAM) (F) is obtained.




Figure 2.4: Level of details with regard to communication and time. Source: [Gajski&09]

Figure 2.5 shows a similar model classification, provided in [Black&04]. Following this classification, BCAM models are called Bus Functional Models and CAM models are referred as Register Transfer Level (RTL) models. Besides different terminology, the interesting issue is that Black shows intermediate refinements both in the communication and functionality aspects that were not provided by Gajski. These refinements fall under the category of Transaction Level Models. Note that, according to this criteria, it is not required to incorporate an approximate notion of time regarding functionality or communication to reach an RTL (CCAM or BCAM) level of representation in one of the areas. However, a complete RTL representation of the system is reached when both communication and computation aspects have been described in the model.




Figure 2.5: Level of details with regard to communication and time. Source: [Black&04]

2.5.4 Outcomes of Hardware/Software co-design

Applying Hw/Sw Co-design in system development provides advantages that are more complicated to obtain from other approaches. Among them, the following must be highlighted:

Enabling technologies

- Development of an architecture fitted for a particular behaviour, optimized according to concrete design criterion.
- Wider exploration of the Design Space, evaluation of multiple architectures before selecting. [Schaumont10] refers to this as "architectural awareness".
- Reuse of IP blocks in order to reduce implementation effort. Pursues to eliminate the *Not Invented Here* syndrome.
- Parallel and cooperative development of hardware and software.
- Consideration of faults, errors and possible problems earlier in the development process.
- Faster time to market, meeting the time requirements of today's business.
- Higher confidence in the achieved design, gained through a model based, well-structured process.

Current research and successful cases demonstrate that a methodology that is making use of Hardware/Software Co-design ends up creating better and higher quality embedded systems. However, one has to consider that implementing such a process in a development team, used to work by following an informal development process is not an easy task. Such a change should be introduced gradually and, to start with, being applied in a small project (risks mitigated). However, methodology adoption considerations are out of the scope of this work.

2.5.5 Enabling technologies

The main technology behind hardware/software co-design is reconfigurable hardware, which is supported primarily by FPGAs. FPGA manufacturers are creating custom toolchains especially adapted for the FPGA hardware they are producing. Both FPGA hardware and development environments are considered as enabling technologies for hardware/software co-design.

2.5.5.1 FPGAs and Hardware/Software Co-design

A FPGA is a semiconductor based device composed by Clustered Logical Blocks (CLB). The CLBs implement the logic present in the FPGA and are communicated through buses. The connections between the CLBs are established depending on the functionality that has to be implemented. Therefore, the wire segments that are conforming the buses are connected to what is called *Programmable Switches*, which state can be modified to fit the deployed implementation. The most relevant feature of the technology described above, is that it enables the system developer to define and reconfigure a hardware structure to fit system design necessities. Some FPGAs allow hardware reconfiguration even in runtime. Such a flexible platform is idoneous for moving a certain functionality from the software the hardware side and vice versa. [Wolf03] refers to FPGAs and its connection with Hardware/Software co-design as follows:

"The platform FPGA seems to be the chip for which co-synthesis was created: The chip's internal architecture is exactly what hardware/software partitioning targets."

Current researchers are using FPGAs as primary target platform for the application of Hardware/Software co-design techniques. Material costs are decreasing considerably and its versatility, among other exotic features, is making them the perfect technological choice for many designs. Acceptance in industry is in constant increment.

2.5.5.2 Hardware/Software development environments

FPGA vendors are providing the end-users with development environments that facilitate the creation of hardware and software blocks. Taking as example the tools offered by Xilinx and Altera, the two main FPGA manufacturers in the market, the kind of tools that are offered are: VHDL Integrated Development Environments, system builders, software IDEs and some auxiliary tools⁸. The used tools in this work are:

System builder: gives the developer the possibility of creating a system with IP cores, processors and its own hardware components developed with the VHDL IDE. The possibility of creating heterogeneous FPGA solutions is supported by this tool.

Software IDE: allows the developer to create software for the hardware system defined in the System builder tool. The application logic developed in this IDE interfaces the hardware side through a Board Support Package (BSP). The BSP contains the hardware dependant software necessary to work with the created custom platform. It can be automatically generated by the tools.

Third parties are providing simulation and modelling tools. An insight on these technologies, and its connection with the hardware/software co-design process, will be given in chapter 3.

2.5.5.3 Other platforms

There are other platforms in the market supporting the partial reconfiguration of hardware blocks. It must be remarked that, even though they present some conceptual similarities to FPGAs, none of them reach the same level of flexibility. In decreasing order of complexity and affordability, some of the current available platforms are: Application-Specific Instruction-Set Processors, Programmable System-on-Chip (PSoC) and microcontrollers featuring Configurable Logic Cells (CLC). Additional information about PSoCs and microcontrollers with a CLC can be found in www.cypress.com and [PIC10F320] respectively.

2.6. Summary

Due to the complexity in nowadays system, abstractions and models making use of them have emerged as a necessity. The chapter has introduced the concept of abstraction and applied it in the context of system design. Different abstraction levels in hardware/software systems have been presented, showing which details can be *described* and *analyzed* at each of them. System-Level Design has been presented, remarking its connection with modelling languages and technologies. Finally, a certain way of doing System-Level Design, called Hardware-Software Co-design, has been described. The partitioning problem, aspects under study and role of modelling in this approach has been introduced.

In order to reach a better understanding on how modelling technologies are used in the development of new hardware/software systems, it is needed to conduct a more thorough review of the existing languages and methodologies in the market.

⁸Front-ends to configure additional hardware or tools to support the debugging process. For example, the application Chipscope, from Xilinx, makes it possible to insert a logic analyzer block in the FPGA and monitor the signals in different lines.

Modelling techniques, analysis and methodologies in hardware/software co-design

This chapter provides an overview of the different modelling technologies and methodologies used for the hardware/software co-design of embedded systems. We start with the introduction in section 3.1, discussing the necessity of abstract modelling languages for representation and analysis. Section 3.2 reviews different modelling languages used for hardware/software design. Time representations used by the languages VDM-RT, SystemC and VHDL are explained in section 3.3. Section 3.4 analyses some methodologies using the modelling languages described in section 3.2. Finally the chapter is summarized in section 3.5.

3.1. Introduction

Multiple authors agree that there is a need for platform-independent languages, not directly associated with hardware or software implementations [Huang&04, Gajski&09]. Platform independent languages do not alter the partitioning process and avoid biasing. The consequence of avoiding biasing is to perform partitioning based on technical reasons and not based on the modelling technology used. Wolf in [Wolf03] refers to the modelling languages and their impact on partitioning decisions as follows:

"Software languages like C bias the implementation in favour of software, while hardware languages bias the results towards hardware."

As an alternative, it is proposed to describe the system by applying an *heterogeneous* modelling approach. This approach proposes to:

- Use a hardware modelling language to represent behaviour that has to be *clearly* allocated in hardware and
- use a software modelling language to represent the rest of the system.

Note that, if it can be foreseen that certain part of the system clearly has to be implemented in hardware, there is no partitioning challenge in that case. Note as well that, if [Wolf03] approach is

followed, the rest of the system, considered as "not-clearly hardware", will be modelled by using a software description language. According to the same article "Software languages bias the implementation in favour of software". It will be fair to conclude that this approach will be biasing the whole implementation towards a software solution.

Even though Wolf's reasoning is too simplistic, the connection between the modelling language used and the final design decisions is well-founded. An alternative approach is to make use of more abstract modelling languages, that are not clearly connected with a particular implementation technology. The languages Parallel Object Oriented Specification Language (POOSL) and the Vienna Development Method (VDM) notations are examples of abstract modelling languages. Both languages are not connected with particular implementation technologies and can therefore be considered as platform-independent.

There are a number of authors proposing to apply heterogeneous modelling, using abstract modelling languages, different Models of Computation (MoC) and graphical modelling languages ([Gajski&09, Shaout&09, Patel&04]). The languages System Modelling Language (SysML) and Unified Modelling Language (UML) are examples of graphical modelling languages. The MoC used for System Level Design (SLD) are out of the scope of this thesis work so they are not considered further in this thesis (additional details are provided in [Gajski&09]).

This chapter provides a description of the modelling languages presented above and some of their associated methodologies for SLD. Additionally, the specialized languages MATLAB/Simulink and SystemC are presented. Even though they are not abstract languages, the fact that they are widely applied in the hardware/software system development field is making them worthwhile consideration.

3.2. Modelling languages

Two modelling approaches are considered in this work: a) informal modelling and b) formal modelling. These approaches could be defined as follows:

Formal modelling: that makes use of rigorous techniques, based on discrete mathematics in order to create models. These models are strong in the use of abstraction and precision. Formal modelling can be applied in the specification, design and verification of software and/or hardware systems. An example of a formal modelling language could be VDM.

Informal modelling: which does not make use of mathematical techniques in the process of model creation. Even though they lack formality and, to some extent, precision, they are strong in communication. An example of an informal modelling language could be the Unified Modelling Language (UML).

Current approaches in industry are making use of formal modelling mostly in the verification stage of the project, as a way of ensuring that the system hold certain properties [Gajski&09]. This section is making a distinction between formal and informal methods, but it is important to keep in mind that the differences between the techniques can be exploited in order to get more communicative models at variable levels of rigour [Bruel&98].

3.2.1 SysML

The System Modelling Language (SysML) is a graphical modelling language developed specifically to support Model-Based Systems Engineering (MBSE). SysML is a UML profile that is

SysML

reusing some of the UML semantics and incorporating additional ones. This relation is depicted in figure 3.1.




Figure 3.1: Venn diagram depicting SysML/UML relationship. Source: [SysML1.2]

According to [Rosenberg&10], the four pillars of SysML are supporting modelling the aspects of the system in terms of *requirements*, *structure*, *behaviour* and *parametrics*. The application of requirements diagrams is not directly connected with this thesis work, however the rest deserve additional consideration:

Structure: represented in SysML by two kind of diagrams: a) *block diagrams* and b) *internal block diagrams*. Block diagrams can represent hardware, software or firmware blocks. Internal block diagrams are used to illustrate the internal composition of each block. They are making use of parts, ports and connectors.

Behaviour: expressed in SysML by four kind of diagrams: a) *use cases*, b) *activity diagrams*, c) *sequence diagrams* and d) *state machines*. Note that these elements are inherited from the UML language, which only uses them in a software representation context. On the other hand, SysML make use of them in order to illustrate behaviour in heterogeneous systems.

Parametrics: enables the incorporation of concrete and detailed characteristics and constraints that the system might present. A *parametrics diagram* makes use of constraints blocks in a certain analysis context.

If the elements presented above are analysed, they could be classified in two groups:

- (a) New graphical representations: Under this category the diagrams for representing requirements, structure and behaviour can be found.
- (b) New representations that enables additional analysis: Parametrics makes possible to integrate analysis with graphical design models.

It can be conclude that: new graphical representations enable the application of a UML based approach in a generic context, make it possible to incorporate a wider range of details (hardware and software) and improve the communication outcome of the models; new representations that enables additional analysis, parametrics, is one of the most relevant SysML contributions, since it enables the evaluation of engineering trade-offs. As it was explained in section 2.5.1, trade-offs

are widely present in the Hardware/Software Co-design field.

Even though, SysML is not a methodology and is tool independent, tool support is relevant in the exploitation of parametrics. A second tool related issue is the *code generation* support. Some tools, like Enterprise Architect or Artisan Studio, enable the code generation in different languages. Here it is important to remark that not only software languages, but hardware description languages are considered as well. Under the appropriate tool support, a SysML model can be used to perform design space exploration, partitioning and co-synthesis of the hardware and software components. A complete case study, illustrating the application of SysML diagrams, usage of parametrics and code generation, can be found in [Rosenberg&10].

- SysML provides intuitive and expressive graphical notation to represent *structure* and *behaviour*.
- *Parametrics* enables the representation of characteristics and constraints by the use of numerical expressions, that can be analysed in constraints blocks diagrams.

3.2.2 SystemC

SystemC is an open source class library for the programming language C++ to model digital hardware. Chen et al. in [Chen&09] show that such a modelling approach is possible due to the integration of hardware constructs and a simulation kernel. [Black&04], elaborate further on the SystemC hardware modelling enabling features, which are summarized below.

Time model: which is managed by the simulation kernel. Time is implemented in kernel space, with a resolution of 64 bits. The time model provides a notion of time to the simulation, but it does not provide the possibility of instantiating clocks.

Hardware data types: that are close to the hardware representation of data. For example, integers and fixed-point representation of a user defined width and the four-state logical data type.

Module hierarchy: so it is possible to represent complex hardware/software structures. Modules can be connected to other modules in the upper, lower or at the same level using channels.

Communication: simple channels and complex communication buses can be modelled. The first type with the integrated communication constructs (for example, fifo queues, signals or ports) the latter, making use of the language in order to model them.

Concurrency model: implemented as a cooperative multi-tasking model. This means that processes are responsible to suspend themselves in order to give processor time to other process. In other words, the simulation kernel *never* initiate a context switch (a non pre-emptive model).

SystemC present a high dynamic range in terms of abstraction [Patel&04]. It is possible to apply it in order to create SAM models and refine it down to the Register Transfer Level (see [Black&04, Gajski&09, Chen&09]). Some tools support the generation of VHDL code from SystemC. This bridges the gap between the hardware and the software development world.

On the other hand, it is generally argued that the use of a programming language in a modelling

context, influences towards the incorporation of implementation details in wrong levels of abstraction. In some cases it is needed to incorporate implementation details, for example, while performing refinement. In other cases it is negative, for example while creating a model that is used as initial specification of the system. The key is, from the modeller perspective; a) to have the abstraction level in which the model under creation is supposed to be clear in mind and b) to incorporate implementation details only when they are necessary.

Finally, remark that [Wolf03] stated that applying a software modelling language to describe a system completely, might bias the design decisions towards software solutions.

The last disadvantage SystemC presents, even though it is not exhibited by the class library itself, is the overwhelming price of the simulation and modelling tools. Even though SystemC can be compiled with the open source compiler *gcc*, and executed with a free implementation of *bash* its capabilities cannot be fully exploited without the appropriate tool support. Due to its possibilities and integration with other technologies, special mention deserve the commercial tools *CoFluent Design*¹ and *Modelsim*². Application of this tools in the co-design process are out of the scope of this thesis work, further details to be found in [Robert&10].

- SystemC is strong in modelling hierarchy and structure.
- The abstraction range covered by this modelling language is very high, ranging from overall system representation to Cycle Accurate Models.

3.2.3 MATLAB - Simulink

MATLAB is an environment that enables the user to perform complex mathematical calculations. It is optimized to offer better performance than ad-hoc models developed in programming languages like C, C++ or FORTRAN [Matlab]. *Simulink* is a MATLAB extension, used for dynamic systems modelling and simulations. Its functionality is organized in toolboxes depending on their application are. These areas range from physical modelling to system verification and validation [Simulink]. Code generation facilities are integrated as part of Simulink, enabling the tool to generate code in the programming languages C, C++ and VHDL among others.

Simulink makes it possible to model particular implementations at the TLM level. Furthermore, hardware blocks provided by Xilinx are already modelled in the tool. This makes it possible to create system models in the form of block diagrams and analyse the performance before having the actual hardware. Finally, the generated model in Simulink can be verified against the final implementation. The implementation can be simulated in a software environment or deployed in hardware. This link is performed by the Simulink component *EDA Simulator Link*.

The main advantage that is offered by the combination of MATLAB-Simulink, is the strong tool support in the modelling of complex systems in terms of: a) Physics and b) Mathematical algorithms. The outcome of using this modelling technology will reach its maximum only if the problem under study presents a challenge in one of those areas (or both).

¹Cofluent official website: <http://www.cofluentdesign.com/>

²Modelsim official website: <http://model.com/>

- MATLAB - Simulink has a remarkable performance when it comes to model mathematical based processes. Special attention deserves its application in the signal processing field.
- There are a number of extensions that allow the application of this modelling technology in a number of fields.
- The language is not appropriate to model complex structures or concurrent problems.

3.2.4 POOSL

POOSL stands for **Parallel Object-Oriented Specification Language**. This language was created at the Technical University of Eindhoven with the purpose of supporting Hardware/Software Engineering through System Level Design (SLD). [POOSL] describes its syntax as expressive and composed by a powerful set of primitives. These syntax primitives are organized around three layers:

Data layer: POOSL is an Object Oriented language, so data entities are represented by classes. Information is represented by attributes and associated operations by methods. Data classes are passive and their methods are sequential.

Process layer: provides the constructs to represent basic active classes. Active classes can be combined to represent composite processes. Process classes give the possibility of expressing *concurrent* and *active* behaviours.

Architecture layer: POOSL gives the possibility of defining architecture by the use of clusters. Clusters are composite structures that combine clusters, process classes and data classes. A cluster does not extend the behaviour of the individual elements composing it, but uses them to offer a more complex functionality. Clusters define their interfaces with external components by the use of messages. Creating clusters inside clusters gives the possibility of representing *hierarchy* in POOSL models.

There is an UML profile for the graphical representation of models created under POOSL. Table 3.1 shows the relation between the layers and the stereotypes introduced by the POOSL UML profile (see [POOSL]). In figure 3.2 the graphical representation of these profiles are shown.

Syntax layer	Equivalent UML stereotype
Data Layer	«data»
Process Layer	«process»
Architecture Layer	«cluster»

Table 3.1: Equivalence between POOSL constructs and UML stereotypes.

It can be concluded that: the general idea about using the POOSL language (formal representation), together with the POOSL UML profile (graphical, informal representation), is a clear application of the ideas proposed by [Brue&98] (previously discussed at the beginning of section 3.2).

VDM-RT

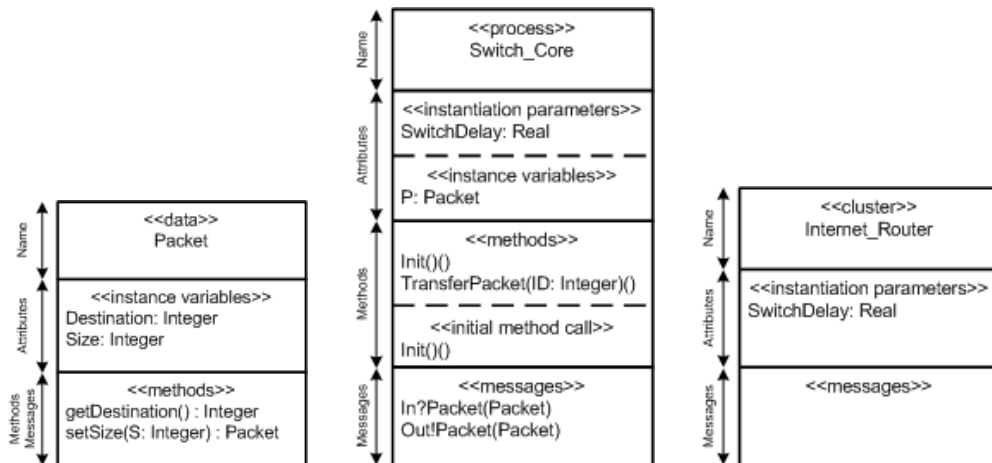


Figure 3.2: Representation of POOSL classes under POOSL UML profile. Source: [POOSL]

POOSL is supported by the tool *SHESim*. SHESim is a graphical tool that allows the interactive composition of the system through blocks. Blocks can be detailed with attributes and methods and other elements depending on their type. An initial graphical model can be refined with the use of the actual language. SHESim can interpret the model so execution can be performed.

Code generation is partially covered by the tools. SHE integrates a formal framework for the *generation of Real-Time Control Software*. However, computation-intensive software and hardware synthesis are not currently available.

POOSL is the most relevant enabling tool of the SHE methodology, which will be described further in the text.

- POOSL presents platform independent semantics.
- A combined graphical-formal approach has been developed for its application in POOSL.
- POOSL allows the creation and execution of formal models in order to verify properties.
- The lack of current research makes it not recommended for new projects.

3.2.5 VDM-RT

VDM-RT is the **Real Time** extension for the modelling language/formal method technique VDM, the **Viena Development Method**. The pillars of VDM rely on the the VDM-SL ISO standard [ISOVDM96short], [Plat&92]. In order to support new problems and satisfy industrial needs, the language has evolved and several dialects have been developed. The considered dialects in this work are:

VDM-SL: is the standardized base language use in the formal specification of systems.

VDM++: is the object oriented extension of the language. Concurrency support by the means of thread and synchronization predicates is provided.

VDM-RT: is the real time extension of VDM, which allows the modelling of distributed and real time systems.

Tool support exist on the market in the form of open-source and commercial solutions. Some of the available tools are:

VDM Tools: commercial tool supporting the base language VDM-SL and the dialect VDM++. It supports the code generation in the programming languages Java and C++. Reverse engineering can be applied over Java code. It is possible to generate UML diagrams from the VDM model and perform test coverage analysis [VDMTools].

Overture: open-source tool supporting the base language VDM-SL and the dialects VDM++ and VDM-RT [Larsen&10c].

One of the most relevant characteristics of the VDM-RT language is the possibility of exploring different deployment options. From the model perspective it allows to perform static system deployment. This deployment is made through the use of the CPU and BUS classes.

CPU: represents a single processing unit. VDM modelled functionality can be allocated on a CPU which will execute it. CPU operation can be configured to run under two different models First Come First Served (FCFS) or Fixed Priority (FP). Under the FCFS mode, the CPU will execute a certain number of operations per thread before it is swapped out. On the other hand, under the FP mode, the CPU will execute threads depending on its execution priority . Besides the operating policy, the frequency at which the CPU is running has to be set.

BUS: Buses serve as communication channels between CPUs. Each bus has to be configured in terms of transmission policy, transmission speed and connected CPUs. The available transmission policies at the time of writing are FCFS and Carrier Sense Multiple Access/Collision Detection (CSMA/CD). In an static deployment context, buses are configured before the model execution start, and remain in the specified topology until the end of the model execution.

Additional work has been carried out in the field of dynamic reconfiguration of VDM models [Nielsen10a]. The application of this technique would allow the modelling, simulation and analysis of changing network topologies in a VDM setting. Such an approach would be worthwhile considering in the modelling of reconfigurable computing systems, which is a hot research topic with potential applications in the FPGA technology. However, this field is out of the scope of this thesis.

Additional work has been carried out in the application of these dialects in concrete industrial cases. Special consideration deserves the application of VDM-RT in modelling and validation of distributed embedded real-time control systems [Verhoef05, Verhoef08].

For additional information on the VDM modelling language in general and on VDM-RT in particular, refer to [Fitzgerald&05] and [Larsen&10b] respectively.

- VDM allows the creation and execution of formal models in order to verify system properties.
- VDM presents platform independent semantics in all its flavours.

Language comparison

	SysML	SystemC	Simulink	POOSL	VDM-RT
Hierarchy and modularity	++	+++	++	+++	+++
RTOS support			+	++	++
State transition	+++	++	+	++	++
Concurrency and Synchronization	+	+++	+	+++	+++
Timing	+	+	++	++	+++
Constructs to support formal verification		+		+++	++
Team communication possibilities	+++	++	++	++	+

Table 3.2: Modelling language comparison

- Strong in representing hierarchy, structure and behaviour in concurrent and real time scenarios.
- Growing tool support and used in on-going research projects.

3.2.6 Language comparison

Table 3.2 presents a comparison between the presented modelling languages and the different desirable features described in section 2.4. The languages have been assessed with a scale from zero to three "+". Zero "+" means that the discussed feature is not supported, three "+" represents a good support of the considered feature.

In general "Hierarchy & modularity" are properly supported in most modelling languages. The support for "RTOS and timing" analysis is weaker in modelling languages based in graphical notations. Executable languages are stronger in representing Concurrent and timing aspects. Finally, graphical languages are stronger in supporting communication across team members. It could be conclude that a single modelling language cannot cover all the aspects that require representation and analysis in a Hardware/Software project. A modelling approach in which several modelling languages are used in order to study the aspects in which they are stronger is the most promising avenue.

3.3. Time and modelling languages

This section aims to explain how time is represented in different modelling languages and how time information is affecting system simulation and modelling. The notions of time introduced in here can be used in order to make a logical system working on discrete basis, react to a continuous phenomenon like time. Such a system is considered as *Real-Time System* [Burns&09].

VHDL and SystemC make use of a notion of time that can reach the femtosecond level [Perry02, Black&04]. Such a precise notion of time is needed because both languages are used to describe hardware/software systems at the implementation level.

No notion of time is incorporated in the modelling languages VDM-SL and VDM++. In the case the modeller wants to incorporate a representation of time in models created in these languages, it will have to be modelled explicitly. A different approach is presented in the real time version of

the language VDM-RT. The VDM-RT kernel incorporates the notion of time by using *time-ticks*. Each time-tick corresponds to the physical notion of time 1 nanosecond.

Additional explanation deserves the real-time mechanisms incorporated in VDM-RT. Time is considered by new types of threads, and the expressions duration and cycles.

Two different types of threads are available in VDM-RT:

procedural: threads, statements executed until completion. They are subject to scheduling and descheduling.

periodic: threads, offer the possibility of repeated execution of the instructions contained on it. The execution frequency together with additional parameters can be set up at declaration time.

Periodic threads can be precisely adjusted to model the real-time problem under study. These threads can be tuned by the use of the constructs **period**, **jitter**, **delay** and **offset**. This work makes use of the **period** construct, which express the time interval between successive executions of the periodic thread. The inverse of the period is equal to the threads execution frequency.

The VDM-RT language provide constructs to use the notion of time on the created models.

duration: specifies that a certain amount of time units are going to be used in order to execute the accompanied statement. This duration is *absolute* and *platform-independent*. This implies that no matter the speed of the processing unit, the instruction is going to take x time ticks to be executed.

cycles: specifies that a *relative* amount of time units are going to be used in order to execute the accompanied statement. The instruction cycles expresses the number of CPU clock cycles required to complete execution. This time information is *platform-dependant*, total time required for complete model execution will depend on CPU frequency.

VDM-RT simulation is time-driven. The VDM-RT interpreter integrated in Overture is implementing threading under a *pre-emptive* paradigm. The scheduler gives a certain amount of time so each thread can perform computations during this time slice. The duration of this time step is determined by the minimum duration presented by one of threads considered in the simulation. In figure 3.3 it can be seen how the time step duration is determined for a concrete case. In this scenario two CPUs running two different threads are modelled. Since thread 1 running in CPU 1 is taking 2 time units to be executed, that will be the size of the minimum time step. Besides the scheduling time slices considered above, different policies can be used. These policies are established with the constructs explain in the beginning of this section.




Figure 3.3: Threading and step size in VDM. Source: [Lausdahl&10b]

In figure 3.4, the resource scheduling process together with the thread policy section is shown. As it is shown there is a central component named *resource scheduler*, responsible for pointing all the system CPUs at each time step. Each time a CPU is pointed by the resource scheduler, it is pre-empted for execution. Once a CPU is selected a computation time equals to the time step is given to it. During that computation time, a certain policy previously defined for the thread will be applied. In the case the duration of the thread reach an equal value to the assigned time step, the scheduler will swap out the thread and assign computation time to the next CPU. It is possible to make such a simulation because there is no drift in the notion of time of each CPU. In other words, all the CPUs used in the model are synchronized.

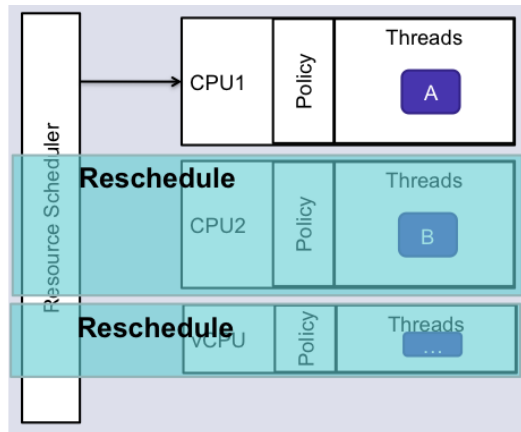


Figure 3.4: Resource scheduling. Source: [Lausdahl&10b]

SystemC and VHDL can represent and conduct analysis at the femtosecond level where as VDM-RT is operating at the nanosecond level. This implies that there is a precision gap of 6 orders of magnitude between the modelling technology VDM-RT and SystemC and VHDL languages.

3.4. Methodologies

The languages introduced in section 3.2, cannot produce results if they are applied in an unsystematic way or isolated. Therefore, they are used within a methodology or development process. There are two ways of using models in the development process: a) use ad-hoc models in order to illustrate and analyse particular aspects of the system and b) use models as a guide in the process of model-driven engineering approach. An example of the first case: the application of MATLAB models in order to develop the signal processing component of a system. The rest of the system is not modelled in MATLAB and it is developed using a different approach. The second approach, model-driven engineering, is applied in all the projects in which a Hardware/Software Co-design methodology is used. Since the latter kind is the core part of this thesis, this section will be primarily focused on this approach.

3.4.1 Xilinx methodology

The non-modelling based methodology introduced in this section is proposed by the tool and FPGA vendor Xilinx. It must be remarked that, this is a tool oriented methodology, and it is not forcing or pursuing the application of any SLD methodology like Hardware/Software Co-design. Therefore, they are mere guidelines on how their technologies and development environments have to be used. This methodology is located at the implementation level, while the rest presented in this section start at a higher level of abstraction, considering the system as a whole.

Figure 3.5 shows the tool flow recommended by Xilinx. This flow is used in the creation of the hardware/software system. As it is depicted, the design starts with hardware. Once the hardware platform is defined, a BSP containing the hardware drivers is generated, so software developers can create the software logic.

It is important to consider this methodology in the appropriate project stage. Considering a project in which Hardware/Software Co-design is applied, several project phases have to be accomplished before using Xilinx methodologies described in [XilinxToolFlow11, Speedway11EDK]. If they are applied before, the development team might face the risk of starting the implementation phase too early, without creating a detailed enough specification or without evaluating enough possible architectures.

The purpose of this section is to show that tool vendors might propose ways of implementing systems, which are highly biased by the tools they are selling. Pragmatism and analysis of vendors methodologies have to be done before incorporating them in the whole development process.

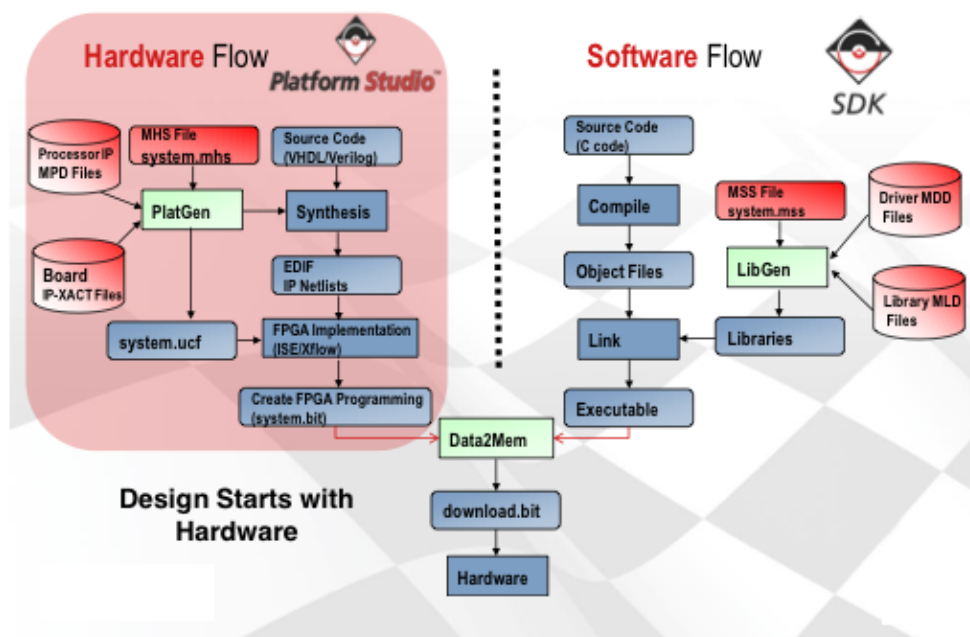


Figure 3.5: Xilinx Microblaze development tool-flow. Source: [Speedway11EDK]

3.4.2 A MATLAB - Simulink based methodology

The official MATLAB methodology is considering primarily the verification of hardware/software designs. A brief summary of the verification approaches supported by MATLAB - Simulink is provided in the appendix E. Since this thesis is focused in the partitioning process, these verification strategies are not relevant. There is research work using MATLAB/Simulink as a tool for

design space exploration. [Ou&05] is proposing the application of MATLAB in order to evaluate different partitionings through the use of model based simulation. This approach is aiming to create high-level abstractions in order to represent processors, buses and dedicated hardware blocks. The equivalence between models and actual implementations is shown in figure 3.6. [Ou&05] is considering these models as high-level abstraction representations, while still they are defined as cycle-accurate. Some authors do not consider, cycle-accurate models to be at a high abstraction level. The explanation behind is that, according to Ou, the simulation will consider the number of actual clock cycles used to perform each computation. Instead of cycle-accurate, the model could be more precisely considered as *cycle aware*. Since not all the cycle related information is taken in consideration, the simulation performance regarding time is improved in orders of magnitude.




Figure 3.6: Proposed approach by [Ou&05]

Note that this modelling approach is not tackling the energy performance of the simulated solutions. According to [Ou&05], additional research has to be done in order to incorporate consumption considerations. One of the proposed ways to model and analyse this issue is by adding an instruction-level energy estimation. Energy estimations require a more detailed modelling of the final implementation [Ou&04] and are out of the scope of this thesis work.

3.4.3 A combined SysML - SystemC based methodology

This approach is proposed by the EU-FP7 project SATURN. SATURN (SysML based modeling, architecture exploration, simulation and synthesis for complex embedded systems). An overall description of the phases proposed by SATURN is presented in [Mueller&10]. The design process starts with modelling the requirements by using SysML *Requirements diagrams*. Structure and behaviour can be described by the use of the correspondent diagrams presented in section 3.2.1. The SATURN modelling frontend, based on *Artisan Studio*, provides tool support to generate SystemC and C code from this graphical description. Generated results can be co-simulated a SystemC simulator is supporting the hardware simulation and QEMU is supporting the software simulation in a virtual machine. Further analysis can be performed with MATLAB/Simulink. Once the co-simulation of different candidate architectures has been performed, and a final one has been chosen. Hardware can be synthesized and software loaded in the FPGA CPU. The infrastructure supporting this process is shown in figure 3.7.

One of the relevant proposals introduced in the SATURN project, shown in [Mueller&10], are the equivalences between the SystemC language and the graphical notation based on SysML. Similar considerations were presented in [Mischkalla&10, Prevostini&07]. In figure 3.8 the applied equivalences in this case are shown.

There is not much documentation about the SATURN project because it is still under development.

3.4.4 The SHE methodology

SHE stands for **S**oftware **H**ardware **E**ngineering methodology. This methodology is using the POOSL modelling language introduced in section 3.2.

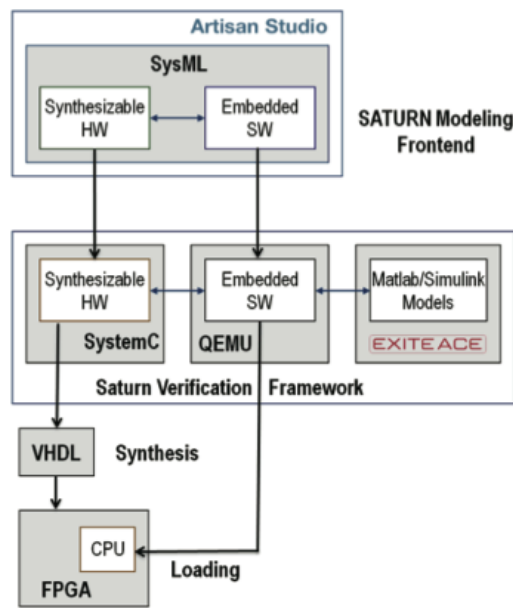


Figure 3.7: SATURN SysML based methodology. Source: [Mueller&10]

SystemC	Stereotype	Metaclass	Notation
Module	<<sc_module>>	Class	
Interface	<<sc_interface>>	Interface	
Port	<<sc_port>>	Port	
Primitive Port	<<sc_in>> <<sc_out>>	Port	
Signal	<<sc_signal>>	Property, Connector	
Process	<<sc_method>> / <<sc_thread>>	Action	
Main	<<sc_main>>	Operation	None
Clock	<<sc_clock>>	Class	

Figure 3.8: Proposed equivalences between SystemC and SysML. Source: [Mueller&10]

This methodology starts with the requirements stated in plain English. The requirements specification output is composed by several documents and artefacts. These elements are informal and its structure might vary. This depends on the requirements engineering techniques applied in each case. Requirements elicitation process and formulation techniques are not considered by the SHE methodology.

The core issues that might influence the design are extracted from these documents, and an *Essential Specification* is produced. This specification presents essential behaviour and essential structure. As a result of this phase a conceptual solution is obtained. Upon this basic specification *refinement* is performed in order to create an *Extended Specification*. The extended specification phase incorporates details enough to perform design decisions and evaluate engineering trade-

offs. Even though it is not shown in figure 3.9 and according to [POOSL], the proposed solution is validated against the initial requirements. In the case this validation, or some part of it, is not successful, it is possible to go back to the *Extended Specification* and refactor part of it. In the case more radical changes are required it is possible to go back to the initial *Essential Specification*. Once the partitioning is definitive and supported by the model analysis, the implementation takes place.

Implementation is partially supported by the tool SHEsim. Implementation of real time control software is possible, but more complex software or hardware VDHL code generation is not supported.

The SHE methodology and the modelling language POOSL has been applied in a number of industrial cases provided by leading companies in the technology sector (Siemens, Alcatel, IBM and Lucent among others). Despite industrial applications and a considerable list of publications, it is not clear from [POOSL] if this research line is active (last related publication in year 2007 and last industrial case in year 2005).

3.4.5 A VDM-RT based methodology

This methodology relies heavily on the application of VDM models in order to support all the project phases and activities. Due to its integration with the project lifecycle, its application is shown embedded in a V-model. The following approach is proposed by [Larsen&10a].

The application starts with the system specification and requirements capture, supported by the use of use-case diagrams (UML, SysML) and VDM-Specification Language. The result of these two models are reflected in a VDM++ sequential model, that helps on modelling a sequential system description. The principal outcomes of such a model is the hierarchical characterization of the system, introduction of the causal ordering of events and an abstract approximate notion of time. A concurrent model is derived in order to show the concurrent aspects of the system. *Threads* and *synchronization* predicates are incorporated to the concurrent VDM++ model.

Once basic structure and concurrent behaviour has been specified, the modeller can create a *real-time* model of the system by using VDM-RT. The accomplishment of real-time deadlines can be studied at this point. As an additional step, *distributed real-time* models can be evaluated. The models explained above, have evolved from an initial specification to a distributed real-time model (finally refined to the implementation). This flow of information is shown in figure 3.10 by green arrows. In the horizontal plane, purple arrows are representing the acceptance activities that are performed at each level, and how the models developed previously can act as a reference for validation.

The application of the VDM technology is under continuous research. Some of the on-going EU FP-7 projects DESTTECS and COMPASS, are extending the language and the tool support in order to cover a wider area. Both projects are relevant in the hardware/software co-design area. DESTTECS allows the co-simulation of physical systems. COMPASS is focused on formalising the interaction between systems, that together conform Systems-of-Systems. Note that physical interaction with the environment and logical interaction with other systems, are factors that might have an impact in partitioning decisions. Further research on the combination of COMPASS and DESTTECS and its application in hardware/software co-design still has to be done. However, that study is out of the scope of this work.

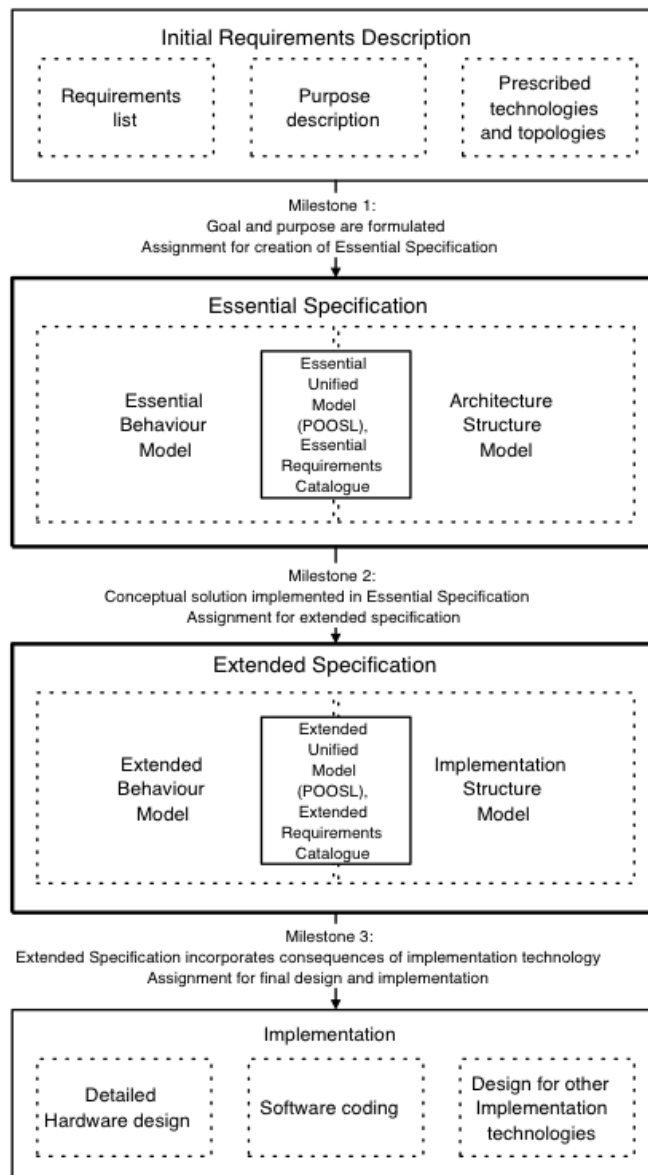


Figure 3.9: Overview of the SHE engineering design methodology. Source: [vanderPutten&07]

3.5. Summary

This chapter has presented several modelling languages with very different approaches to system modelling. Graphical languages like SysML and UML and informal languages like SystemC and Matlab/Simulink have been reviewed. Formal languages like POOSL and VDM have been presented. Different methodologies associated to the reviewed modelling languages have been described. This chapter has shown the need of new modelling languages and/or methodologies that allow the system engineer to put aside implementation considerations and focus on the problem under study at a high level of abstraction.

Summary

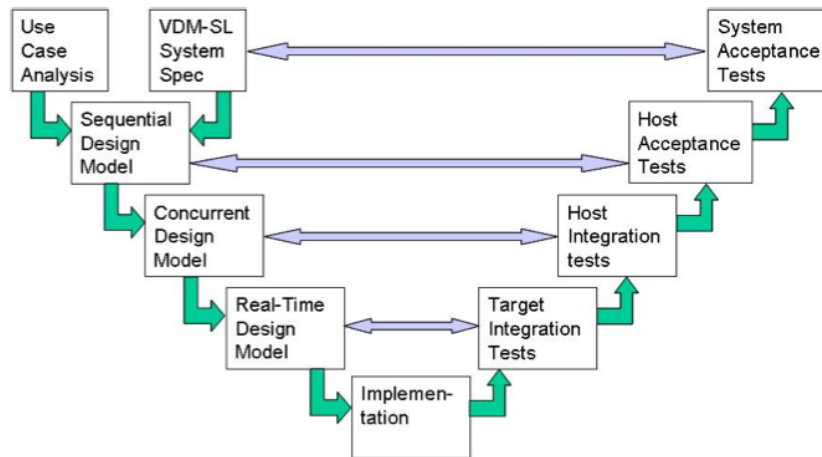


Figure 3.10: VDM based methodology. Source: [Larsen&10a]

A VDM-RT based modelling methodology for Hardware/Software engineering

This chapter is presenting the proposed methodology for the hardware/software co-design of embedded systems, the core part of the thesis. Section 4.1 gives an overview of the rationale and aim of the methodology presented. Section 4.2 describes the proposed methodology. This description is built on top of the existing VDM-RT development process [Larsen&09] and makes special emphasis on how it can be applied in order to take Hardware/Software co-design partitioning decisions. The methodology is complemented by the early model evaluation against a prototype, which is described in section 4.3. Proposed improvements to the VDM technology are presented in section 4.4. Section 4.5 describes how new platforms can be incorporated to VDM-RT. Finally, the chapter is summarized in section 4.6.

4.1. Introduction

This chapter presents the VDM-RT modelling methodology proposed in this thesis work in order to provide support to the Hardware/Software Co-design of embedded systems. The goal is to provide a way to tackle the fundamental problems in hardware/software co-design described in chapter 2. The only avenue to accomplish this objective is to apply a systematic methodology, based on the ones presented in chapter 3. The intent of the proposed VDM-RT methodology is to provide the system engineer with a tool to describe and analyse Hardware/Software systems from a formal perspective, that will allow him to perform design space exploration earlier in the project life-cycle. A number of guidelines based on the approach explained in the chapter will be given. The guidelines give modelling advice according to the hardware/software co-design process proposed in this chapter.

4.2. Methodology description

The presented methodology is structured in three modelling stages and one exploration phase. The methodology starts with the *early* production of sequential VDM models, that will start being far away from the final one but will help on getting familiarized with the problem under study. Finally, a sequential VDM model with enough details to take partitioning decisions later on in the

process will be created. This candidate model will be moved to a concurrent VDM model, with the purpose of focusing on the concurrent behaviour of the system. A real-time VDM model will be derived from the concurrent model, substituting the very limited notion of time for a more precise one. In the exploration phase, the real-time VDM model will be evaluated as a distributed real-time model. It is here where VDM-RT possibilities for supporting hardware/software co-design decisions have to be studied more in detail. Note that this is a "waterfall-like" process, in which errors introduced in the sequential system model can have a very negative impact in the rest of the modelling stages and exploration phases.

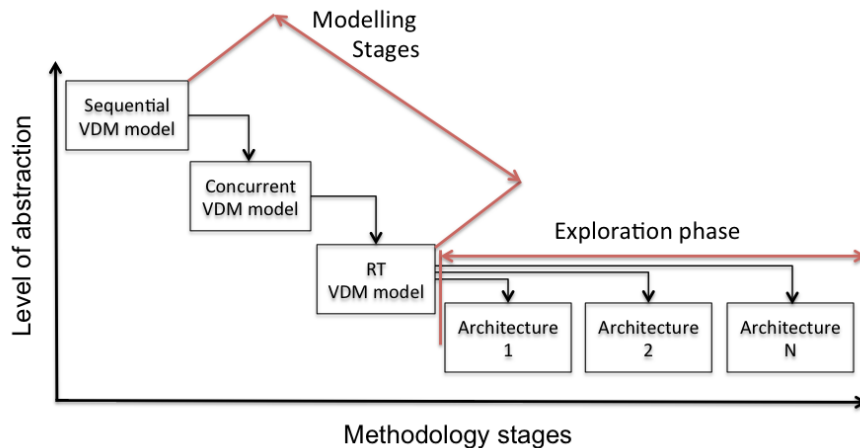


Figure 4.1: Methodology overview.

4.2.1 Creation of a sequential model of the system

The first step of the methodology consist on the creation of a sequential VDM model. As the name says, this model is representing the system operation in a sequential model, modelling the operation execution in a sequence of steps. The notion of time incorporated in this model is very limited, and can only be used to determine whether an event is happening before or after another, considering the system deployed in the same computation unit.

The purpose of such a model is to represent the most relevant entities involved in system operation. The relevant term refer to entities that fall directly under the following categories:

Model entities: represent elements that are a straightforward logical representation of an existing entity in the physical world that has to be processed by the system. For example, a clock class that represents a device to measure time in a model of a time synchronization system (more details about this on a concrete example in chapter 6).

Entities with core behaviours: that represent crucial processes in the system operation. An example of this kind of entity would be an entity responsible for the generation of a Pulse Width Modulated signal controlling a motor (more details about this on a concrete example chapter 5).

System boundaries: that represent the limits of the system. Boundary entities define how and where the system is interacting with the environment.

At this point of time, a "raw" sequential model has been created. Taking advantage of the Object Oriented (OO) techniques, the model can be expressed by a UML class diagram. OO techniques

Creation of a concurrent model of the system

from the software engineering world can be applied, for example design patterns. Once the structure of the system has been created, operational details, used data and system properties can be incorporated.

Guideline 1: The sequential VDM model gives the possibility of focusing on the key hardware/software entities to be considered for partitioning. It is a pre-condition to reach a good understanding of the system before moving to more complex modelling.

Guideline 2: In the case the problem under study is new to the modeller, it is difficult to create a useful model in the first attempt. In this case, the creation of the model should be iterative, and this process should be used as a tool to help on system understanding.

4.2.2 Creation of a concurrent model of the system

The concurrent VDM model of the system is an intermediate step between the sequential VDM model and the real-time VDM model. In this model, the entities with core behaviours most likely will be active classes, with threads running the key operations identified in the sequential model. Additional effort has to be spent on ensuring a safe concurrent access to shared data. The notion of time will remain in discrete timesteps, as defined in the firsts sequential models. In order to make the threads advance progressively, the `TimeStamp` pattern presented in [Lausdahl&10b] can be used.

Guideline 3: It is important to make sure that *all* the synchronization issues are tackled during the concurrent modelling phase. If not, problems may arise during the RT modelling phase. Fixing synchronization problems in the RT modelling phase is altering the focus of that phase and it might be more complex.

4.2.3 Creation of a distributed real-time model

The distributed real-time VDM model introduces three major changes: the introduction of a new notion of time, the introduction of a RTOS layer and the possibility of deploying different model entities in different computation units. This third modelling stage allows to get the models closer to real systems performance.

Guideline 4: The first step in the process of moving to the real-time world is the creation of a non-distributed real-time model. This model is an intermediate step between the concurrent and the distributed real-time model. In this case all the objects will be deployed in the same CPU.

The notion of time introduced by the `TimeStamp` class should be removed and the `Overture` time reference should be used instead. In the case there were threads using `sleep` calls based on `TimeStamp`, they should be refactored to use the information provided by the `time` expression. In the case this sleep period was constant in time, the thread could be modelled as periodic.

4.2.4 Evaluating hardware partitions from VDM-RT models

This methodology proposes the application of VDM-RT CPUs in order to model a hardware partition. It is possible to create a *coarse-grained* model of the hardware partition if this process is carefully applied. The key principle behind it is that there should be a difference of orders of magnitude in the processing speed between CPUs emulating a general-purpose CPU and a CPU emulating an specialized hardware block (hardware partition).

The communication between the CPU emulating the hardware block and the CPU emulating the general-purpose CPU will be modelled as a BUS with a very high bandwidth, orders of magnitude above the bandwidth used in order to represent a pure software-based communication. The reason behind using a high bandwidth is to simulate how a hardware block is using a register mapped scheme for interaction with the control software. This process can be repeated with different alternative system definitions easily. The Overture log files are a powerful tool in order to study the real-time evolution of the system. More details will be provided in the case studies (chapters 6 and 7).

Guideline 5: A VDM-RT CPU that represents a hardware partition should be configured with a processing speed orders of magnitude higher than general-purpose CPUs.

Guideline 6: A VDM-RT BUS that communicates a hardware partition with a general-purpose CPU should be configured with a very high bandwidth. This bandwidth should be orders of magnitude higher than the ones used to communicate general-purpose CPUs running software.

Guideline 7: The instruction **time** will give the value of the notion of time at invocation time. It is a good idea to use it several times during the critical operations under study are being run. This will give an overall idea of the time spent by different operations. Custom logs can be generated, with the advantage of being able to get the relevant data faster than by comprehending the Overture logs.

4.2.4.1 Evaluation of a single hardware partition

The study of a system that is making use of a single specialized hardware partition is a straightforward application of the idea presented above. The steps could be summarized as follows:

1. Identify the component that is going to be executed in hardware.
2. Define a separate CPU configured with a processing speed orders of a magnitude above the speed used in the general-purpose CPUs.
3. Communicate the defined CPU through a bus with a very high bandwidth with the general-purpose CPU. The general-purpose CPU is running the rest of the entities as software components.
4. Deploy the functionality that should be running in hardware in the CPU defined in step 2.
5. Analyse the possible improvement in the system performance.

Guideline 8: In the case the communication to be modelled is not as fast as the one that could be achieved by a registered mapped configuration, the bus bandwidth can be lowered.

4.2.4.2 Evaluation of several hardware partitions

In some situations it is desirable to consider how the system would be performing in the case *several components* are deployed in hardware. Considering the process described above, the most intuitive approach to model such a scenario would be to define a second CPU and deploy the hardware candidates there. However, this process turns to be erroneous and can lead to false results. A more concrete example is presented in order to explain this challenge.

Consider a system that is modelled by three components: A, B and C. Components B and C depend directly from component A, that needs to be communicated with them. Components B and C do not present any kind of interaction between them. The system designer is interested on evaluating the performance of the system when A is running in a general-purpose CPU and B and C are off-loaded to specialized hardware blocks. In the case the process explained above is applied, the system would be deployed as illustrated in figure 4.2.




Figure 4.2: Wrong usage of the hardware partition.

If such a deployment is performed, the specialized hardware block would not be acting as such. In this scenario the CPU representing the specialized hardware block will be behaving as a hi-speed CPU with an OS running two threads. Naturally, such a model could possibly be performing better than a purely software based architecture. These improvements are still not reflecting how a hardware off-loading for components B and C would be performing.

The correct way of performing the hardware off-loading evaluation is to make a deployment like the one presented in figure 4.3. As it can be seen, two CPUs (HardwareA and HardwareB) representing specialized hardware partitions have been introduced. Both CPUs are communicated by hi-speed BUSES to the general-purpose CPU (Controller). Component A is deployed in Controller, components B and C are deployed in HardwareA and HardwareB respectively.

In the case the components B and C would be interacting, a hi-speed bus should be connecting the CPUs in which they are deployed, in this case HardwareA and HardwareB.

Guideline 9: Each hardware candidate should be deployed on an individual CPU configured as a hardware block.

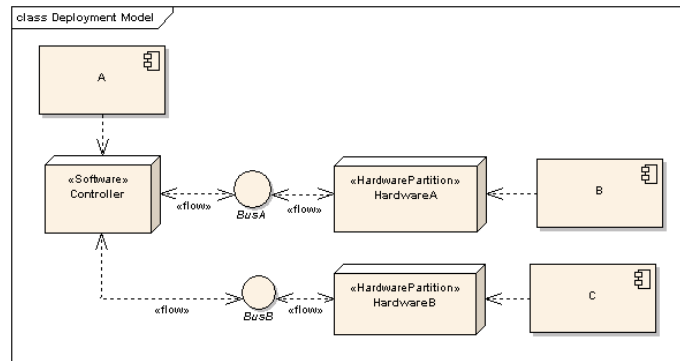


Figure 4.3: Correct usage of the hardware partition.

4.2.4.3 Modelling complex communication between partitions

The BUS class provided by Overture allows the modelling of a bus under the policy *First Come First Served* <**FCFS**>. In some cases it might be interested to make use of other policies. Two possible situations in which this would make sense could be:

1. A system in which not all the information sources have the same priority when it comes to access the transmission medium.
2. A system with different kind of policies depending on if the information is sent or received.

Since these cases are too particular, it is understandable that are not incorporated by default in the language. However, both scenarios can be easily modelled in VDM-RT controlling how the component is accessing the bus. The idea is to make use of a `Proxy` class that implements the desired policy. Figure 4.4 illustrates how this pattern could be applied in the second example presented above. In this case the BUS is communicating using CPU1 and CPU2. The BUS is accessed through a `Proxy` class. This class specifies that a policy has to be applied over the data. However this policy is abstract and has to be implemented in derived classes. `IncomingData` and `OutgoingData` are specific implementations of `Proxy`, that will implement the required policy depending on if the information is received or sent. The interaction with the `Proxy` class within the component is straightforward, in this case there is a controller class `Logic` making use of it.

Guideline 10: Use a proxy class in order to incorporate a more complex behaviour in the transmission/reception of data. As an additional advantage, Proxy classes separate access to the physical communication medium from the system logic, which leads to a lean model structure.

4.2.4.4 Modelling external measuring hardware

During the development of hardware systems it is very common to use external equipment like oscilloscopes or logic analysers. The purpose of such equipment is to monitor the value of a signal over time. There are no classes in VDM-RT implementing oscilloscopes or logic analysers, but certainly some log facilities can be built into the models to provide a similar function. The process to log data is based on a class that obtains the values of the target variable over the model execution time. Once model execution is finished, data is dumped to a file and plotted using a third

Early prototype generation and model validation

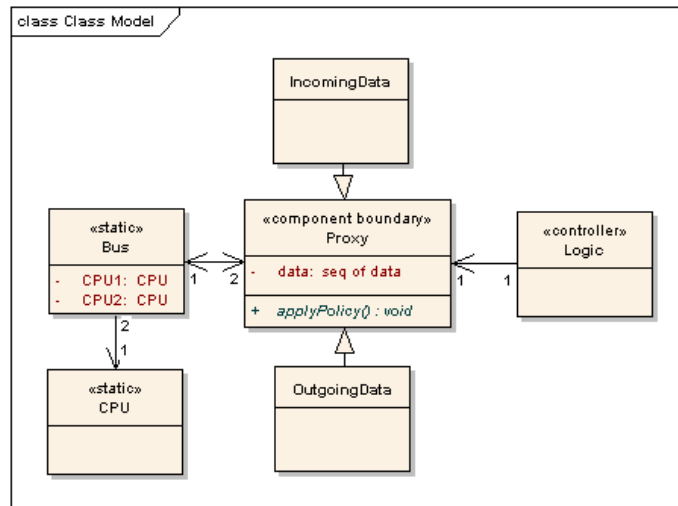


Figure 4.4: Application of a proxy class

party application. The critical part of the logging process is how the data is obtained. Logging implies introducing additional logic in the model. The execution of this logic might alter time related results.

The most simple way to avoid altering time results is to deploy the logging functionality in the virtual CPU and force the class containing the variable to be monitored to store its values periodically in the log class by using a `pushValue` operation. This operation should be defined as an **async** log class method and its invocation should be preceded by a **duration** (0).

Guideline 11: Using graphics is a fast way of evaluating the I/O signals that are fed/generated by a certain architecture. A second advantage is that models generating a visual output are better at communicating modelling results to development team members. Finally, keeping together models and generated plots are a good way to document the designs.

Guideline 12: In general, log operations must be preceded by **duration** (0), otherwise the time spent on logging will affect model execution results.

4.3. Early prototype generation and model validation

Taking one step further in the proposed process, a preliminary prototype implementing some of the proposed functionality could be produced. The value of generating a prototype with regard to the models created in the previous stages, is that it is possible to compare the model predictions with a preliminary implementation. In the case the models have been created correctly, resemblance between the created model and the preliminary implementations is expected, at least in the conceptual model. Some FPGA manufacturers provide tool support in order to profile the produced implementations. Combining the prototype performance, measured by dedicated test equipment (scopes, logic analysers) and assessed with a previously defined criteria with profiling information, can be a way to obtain feedback from a real implementation. In the case the

events/behaviour/performance problems predicted in the models are present in the preliminary hardware implementation, confidence in the generated designs will be supported by well founded design decisions. In the case the obtained results are differing in aspects like time units, while still keeping the expected results, scale factors can be introduced in the models. Finally, if the obtained results are not related to the model predictions, there is a problem in the model, in the prototype, in the measurements or in a combination of all these elements. The solution to such issues depends on the particular case.

Guideline 13: Early feedback incorporation improves the models and bridges the gap between abstract representations and real implementations. In order to exploit the application of models in the development process, it is important to keep aligned models and produced/generated components.

4.4. Proposed improvements to the VDM technology

Since the Overture platform is in continuous development, and this work is carried out in the same institution in which this Master's thesis work is presented, it is interesting to discuss possible additions that could be done in order to improve the tool. Most of these additions/new features are directly connected with the hardware/software co-design field. Others would be beneficial in general, and would increase the value of the tool.

Models of new platforms: It would be desirable to have new specialized hardware blocks ready to use in the form of a VDM-RT library. These blocks could be modelling specialized processors like DSPs, microcontrollers or bus interfaces. A certain set of benchmarks could be run over different platforms and its results incorporated as concrete CPUs able to host components. The benchmarks could evaluate the time it takes to perform certain mathematical operations, communication through built-in buses, etc... Results could be mapped to concrete VDM-RT CPUs. Performance evaluation over these platforms could be more reliable and closer to reality.

Model profiler: Some of the coverage information generated by Overture could be extended so it incorporates the time it has taken to execute each operation. This information can be already extracted from the generated overture log files. However, it is something that has to be done manually for each operation. The possibility of generating such a model profile would simplify the analysis activities, by making the information easier to be read and found.

Improved bus support: The current bus configuration is limited to the First Come First Served policy. A general-purpose bus with a transmission/reception logic that could be user defined would be very interesting from the system design perspective. Such a field could contain the logic to filter packets, introduce noise, establish priorities or cause delays.

Automatic variable monitoring: In some cases it is interesting to monitor the evolution of a variable over time. Overture could incorporate the possibility of automatic plotting of certain variables selected by the user. This new feature is currently under development.

Code generation: Such a feature would bridge the generated models during the design phase with the implementation activities. In the case full code generation is not possible it would

be interesting to have at least the stub of the application. Even though this would not be a considerable leap forward regarding implementation time, it establishes a relationship between models and code. This relation can be backtracked and models can act as a reference for the developers.

Annotations for Real-Time deadline evaluation: A strong point of VDM-RT is the possibility of real-time deadline accomplishment evaluation. In the case a language could be used in order to define the real-time languages, Overture could automatically generate a report after model execution presenting which deadlines have been missed and when. This idea has been evaluated in [Ribeiro&11] and it is currently under development.

4.5. Modelling new platforms under VDM-RT

Some of the proposed improvements to the VDM technology like "Improved bus support" or "Models of new platforms" would require to study the aspects that characterize real BUSES and CPUs, which aspects are relevant for incorporation and which additions to the VDM language would be needed. These additions could be implemented as new constructs or as annotations. In this section some of the required additions to model new CPUs are considered.

Incorporating models of new platforms could be carried out by subclassing the VDM-RT **CPU** class. A microcontroller is an interesting platform in the hardware/software development area. In this case, a worthwhile considering feature would be the addition of interrupts. This would imply that the microcontroller CPU constructor would need to register an interrupt, a microcontroller class method would have to associate this interrupt with a triggering event and, finally, a microcontroller class method would have to associate the interrupt with an operation defined in the logic running in the microcontroller. Finally, the Microcontroller class constructor would have to skip the scheduling policy. This detailed example is shown in the VDM listing below.

```

1 instance variables
2   int1    : Interrupt := new Interrupt(priority);
3   event1  : Event     := new Event(Timer`onTimerOverflow);
4   mcu     : Mic       := new Mic(speed, int1);
5
6   public static timer      : Timer      := new Timer(10,event1);
7   public static counter    : Counter    := new Counter(timer);
8   public static ledControl : LEDdriver := new LEDdriver(counter);
9
10 operations
11 public System : () ==> System
12 System () ==
13 (
14   mcu.event(int1, event1);
15   mcu.setOp(int1,LEDdriver`turnOnLED);
16
17   mcu.deploy(ledController1);
18   mcu.deploy(counter);
19   mcu.deploy(timer);
20 );
```

Considering as a second example a Digital Signal Processor¹, some of the operations deployed on it are executed with the same performance if compared to a general-purpose CPU. However, other mathematically complex operations run orders of magnitude faster under DSPs. Since a DSP is a specialized CPU, subclassing could be used in order to inherit the common features with CPUs. The use of annotations would allow to characterize operations and functions in VDM as DSP efficient without having to alter the description of their logic. This example is shown in the VDM listing below.

```

1  -- DSPefficient
2  private fft: () ==> ()
3  fft() =
4    is not yet specified
5  post doneFlag = true and overflow = false;

```

4.6. Summary

This chapter has presented a VDM-RT based methodology for the Hardware/Software co-design of embedded systems. An overall view of the methodology has been provided and special emphasis has been made on the proposed additions to support the partitioning decisions. A set of guidelines have been presented.

As it has been argued above, there is no perfect methodology in order to design systems. It is the belief of the author that the proposed methodology compensate several shortcomings of existing ones, such as the high degree of specialization of Matlab based workflows, the complex setup of the SATURN methodology or the lack of planning and analysis in a Xilinx based approach. Comparisons aside, *pragmatism* should be present when selecting the modelling technology.

The following chapters 6 and 7 will present two case studies in which the VDM-RT hardware/software co-design methodology is used. These case studies will give a practical insight on the methodology.

¹Specialized processor with a high performance in heavy mathematical computations (e.g. computer vision algorithms, digital filtering algorithms)

Introductory case study: RC servo control

This chapter applies the methodology presented in chapter 4 in a concrete case study. This case study ties together all the previous information presented so far. A model-driven approach will be applied in this problem, and it will be shown how implementation details are abstracted, as described in chapters 2 and 3. This chapter shows the importance of Hardware Software Co-Design. The chapter is structured in four sections. An introduction to basic servo control will be given in section 5.1. The application of the modelling methodology will be presented in section 5.2. An actual implementation will be described in section 5.3. Analysis of the different implementations and reflections on the correlation between the modelled architectures and the final implementations will be provided in section 5.4. Finally, a summary of the chapter is given in section 5.5.

5.1. Introduction to basic servo control

The problem presented in this chapter is the control of a servo motor. Servo motors are electro-mechanical devices able to orientate its axis between 0 and 180 degrees. Servos contain a direct current (DC) motor, several gears, a potentiometer and an electronic control board. The electronic control board acts as an internal servo motor controller and as an interface between the servo and the external controlling device. A detailed internal view of the servo is shown in figure 5.1.



Figure 5.1: Servo internals. Source: hitecrcd.com

The servo electrical connections are: *Vcc* in order to power the servo, *GND* electrical ground required in all electrical devices and *Signal* the control signal used to command the servo. The

control signal is *Pulse Width Modulated*. A signal of these characteristics will change the duration of the pulse according to target position for the servo. Therefore, there is a direct connection between the pulse duration and the final servo position. Note that the frequency of the signal remains constant, the only parameter that changes is the *duty cycle*.

The required command signals vary from one servo to another, and it is something that the servo manufacturer might change. It is common to use pulses between 1ms to 2ms in order to control the servo. The period of the control signal normally varies from 18ms to 25ms (55 to 40 Hz). The period of the control signal depends as well on the servo manufacturer. In order to reach a correct servo operation it is a requirement that the servo is pulsed at the required control frequency, otherwise *jitter* will take place.

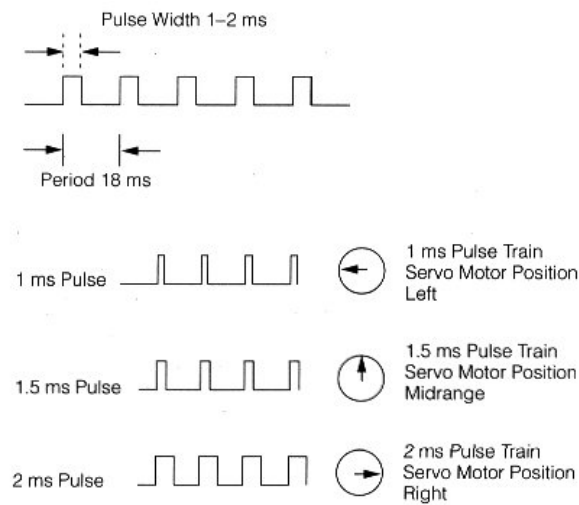


Figure 5.2: Typical servo control signals. Source: seattlerobotics.org

Figure 5.2 is an example showing the working principles introduced above. In this particular servo case the pulse duration is 18ms. This means that the signal control frequency is working at 55Hz. The pulse width will vary between 1ms to 2ms, representing 0 and 180 degrees respectively. Intermediate pulse widths can be used in order to position the servo at different angles of the described semicircle.

The *partitioning decision* in this case is that a servo controller can be implemented in software or in hardware. In the case the servo is controlled from a software component deployed in a microcontroller without an OS, the solution will consist on toggling a logical output between low and high values respecting the required timing constraints. These can be established by inserting properly wait routines in between the toggling functions. In the case the solution is implemented in hardware, the functionality will be exactly the same but implemented in VHDL, and running independently from the software side. Additional details on both implementations will be presented in the following sections. Even though functionality can be reached in several ways, timing constraints will play a major role on deciding upon a concrete implementation.

5.2. Modelling of a servo controller

Following the thesis aim of illustrating model-based engineering using the VDM method, the first step in this case study is to model the problem. The modelling flow will proceed according to the

process explained in chapter 4.

5.2.1 Sequential model

The first step is to create a sequential model of the problem under study. This model will allow to represent unambiguously the entities involved in the system and the causal ordering of events (application of guideline 1).

This model is composed by the following entity classes:

Controller: represents the main control logic responsible for the system operation. In terms of hierarchy this would be the top-module of the system.

PWMgenerator: represents the logic used for the generation of the PWM signal described in the previous section.

Clock: periodic signal regulating the execution of the model. This signal is a logical notion of time: the simulated time.

GPIOinterface: represents a hardware block. It represents the boundary between the microcontroller unit and the environment. GPIO stands for General Purpose Input Output. Each GPIO channel has an associated pin in the microcontroller package. This pin can be at low level (0 Volts) or high level (5 Volts¹). The GPIO hardware block is mapped to a control register that can be loaded with the control value, determined by the algorithm running in the microcontroller CPU.

Additional classes have been used to run the simulation. These classes are mere helpers and not as relevant as the ones introduced above. However they are relevant in order to generate the simulation results.

PositionLogger: keeps track on the target positions issued by the system control logic. This class will generate a Comma Separated Value (CSV) file with the targeted positions.

Sampler: offers the same functionality as an oscilloscope. In this model implementation this class is sampling the outputted GPIO values. It can be attached to different GPIO lines in order to simulate a multi-channel scope. However, this functionality has not been used in the servo case. Once simulation is completed the sampler class is able to dump all the readings in a CSV file.

Figure 5.3 shows a UML class diagram illustrating the relations between the entities presented above. The model is started and configured by the `Environment` class, which will be responsible as well for tearing down the execution (in this case only dumping the acquired signals to files). The `Environment` class will access as well the simulation clock represented by the `Clock` entity. Each time a clock tick is generated by the `Clock`, a sequence of `timeStep` operations will be triggered from the environment in all the instantiated classes. This can be seen in the code listing below. Each class that holds an active responsibility is implementing a `timeStep` operation.

```

1  while (clk.getSimTime() <= simTimeMax) do
2  (
3  currentSimTime := clk.getSimTime();
4  controller.timeStep(currentSimTime);

```

¹Considering Transistor Transistor Logic (TTL)

```

5 | pwmUnit.timeStep();
6 | sampler1.sample();
7 | clk.tick();
8 | );

```

In the case of the PWMgenerator entity, the timeStep operation is triggering the execution of the generateCycle operation. This operation is responsible for creating the control signal that is provided to the servo. The code listing below shows how it has been modelled in VDM. An internal tickCounter variable is representing the number of ticks that will be used to generated a full PWM cycle (composed by the high and low state). The high state of the PWM signal starts at logical tick 0 and ends at logical tick tickLimit. As described in the model, these boundaries will be triggering the setOn and setOff operations causing, therefore, the proper logic level transition in the associated GPIO class instance.

```

1 | private generateCycle: () ==> ()
2 | generateCycle() ==
3 |   cases tickCounter:
4 |     (0) -> setOn(),
5 |     (tickLimit) -> setOff(),
6 |     others -> skip
7 |   end;

```

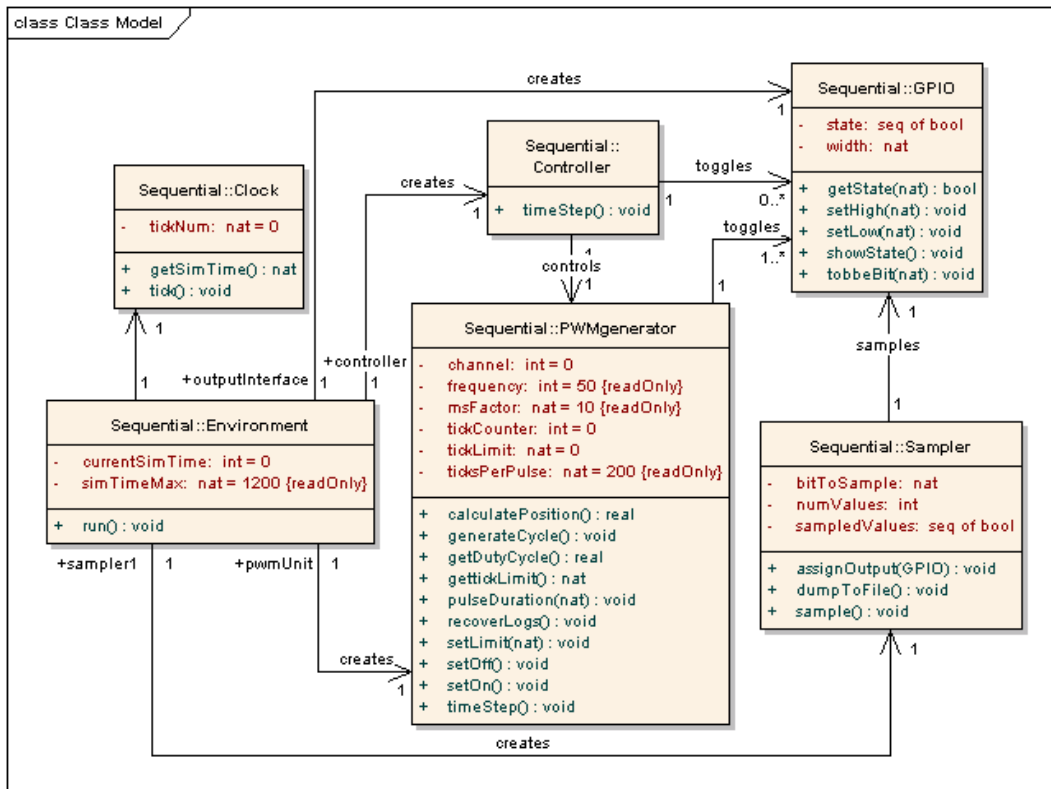


Figure 5.3: UML class diagram view of part of the sequential model.

Concurrent model

The main outcome of the creation of a sequential model of the system is the unambiguous representation of the system hierarchy, and the separation of responsibilities among entities. An additional advantage derived from the application of the VDM methodology is the possibility of executing the models under different scenarios, something that is not that easy to do under UML². The notion of time incorporated in the model is used in order to determine the ordering of events. However this notion of time is not sufficient in order to study time dependant properties of the system, something that will have to be studied in the real time models. Concurrent operation is not incorporated in the sequential model either.

The execution of the model produces as a result a CSV file that can be plotted (application of Guidelines 11 and 12). The resulting plot is a representation of the generated PWM signal by the modelled system. In this case the tool *gnuplot* and a set of bash scripts have been used in order to generate the signal plot. Figure 5.4 show the predicted signal output by the VDM sequential model.




Figure 5.4: Generated servo control signal.

5.2.2 Concurrent model

The concurrent model gives the possibility of showing the concurrent aspects of the systems. The outcomes of such a model are to present behaviours that should be running in parallel and to study how these behaviours are interacting in a safe manner. In this model the timestamp pattern is used. This pattern allows a smooth conversion from the sequential to the concurrent world, by the introduction of a thread barrier. This barrier allows the regulation of the time progression, making sure that at least all the threads have been executed `timeStep` once before making the simulation clock advance one time unit. More details on this pattern can be found in [Lausdahl&10b].

Figure 5.5 presents a simplified UML class diagram of the concurrent VDM model. Attributes and methods are not shown in this class diagram in order to increase readability. These details can be consulted on the previous class diagram 5.3. Stereotypes `«thread»` and `«synchronized»` are used in order to display the concurrent behaviour of the entities.

²Some Computer Aided Software Engineering (CASE) tools available in the market provides the necessary facilities to execute UML sequential models and generate different execution traces. An example of such a tool is IBM Rational Rhapsody. These tools are complex and expensive. An additional advantage is that, executable UML models have to be following the syntax specified by the UML standard. The result is a syntactically correct and executable model.

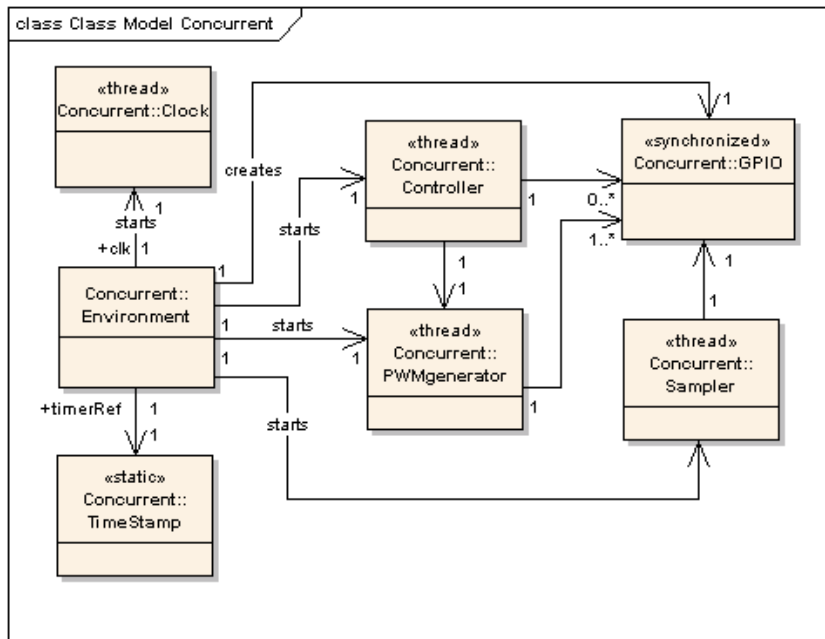


Figure 5.5: Concurrent model simplified UML class diagram.

From the VDM perspective, the model infrastructure is constructed by the `Environment` class. The threads are started from this entity as well, which will block until the finishing condition has been reached.

```

1  start (clk);
2  start (pwmUnit);
3  start (controller);
4  start (sampler1);

```

Then `Environment` class makes sure all the threads have progressed at least one time unit by invoking the static operation contained in the `TimeStamp` class `WaitRelative`. The `TimeStamp` class maintains as well the notion of time used in the simulation.

```

1  while timerRef.GetTime() <= simTimeMax do
2  (
3    timeStep();
4    Environment `timerRef.WaitRelative(1);
5  );

```

The most outstanding change in the VDM models is that now all the active classes, which are running in threads, are incorporating a thread section in the class body. This thread section defines the concurrent behaviour the class is going to present. Thanks to organizing the model evolution in the form of discrete `timeSteps` contained in a single function, the evolution towards a concurrent version is relatively simple. As an example, the concurrent aspect of the `PWmgenerator` class is presented in the code listing below. As it is shown in the model, the `timeStep` operation will be running once and then the thread will block until it is notified by the `timeStamp` instance

class (declared as `timerRef`). The thread will be executing continuously until the Environment class finishes the execution.

```

1 thread
2   while true do
3     (
4       timeStep();
5       Environment `timerRef.WaitRelative(1);
6     );

```

The `timeStep` operation invoked from the while loop, is modelling exactly the same behaviour as the one presented in the sequential model subsection. Besides detailing the active behaviour of the system, synchronization primitives have been used in order to ensure a safe concurrent execution (application of guideline 3). The primitive **mutex** has been used in order to express the necessity of mutual exclusion between readings and writings during the information exchange. The most significant case is present at the GPIO class, which is updated by the `PWMgenerator` and sampled by the `Sampler` class. The used protection is shown in the next code listing. As it is modelled, the state presented by a certain gpio bit can be consulted if it is not being updated. Additionally, a certain bit cannot be being altered at the same point of time.

```

1 mutex (setLow, setHigh);
2 mutex (setHigh, getState);
3 mutex (setLow, getState);

```

At this point of the modelling stage, a model that presents proper structure, hierarchy and behaviour, besides a concurrent operation description has been reached. The next step in order to study how the system behaves in real time and under different deployments (different architectures).

5.2.3 Distributed real-time model

The distributed real time VDM model brings the possibility of exploring the real time behaviour of the previously developed models. Besides studying real time properties, different architectural candidates can be evaluated. In order to model these architectures, the `deployment` class has been incorporated to the model:

Deployment: defines the system infrastructure by the use of processing blocks (CPU) and communication interfaces (Buses). It contains the static class declarations of the modelled entities. This class is also responsible for deploying the modelled entities in the processing units.

Since the notion of time is provided by the overture kernel, the clock class is not needed. The changes in the classes can be seen in the class diagram depicted in figure 5.6. Methods and attributes are not displayed in order to increase readability of the diagram. For additional details refer to figure 5.3.

Initially, a non-distributed real time mode has been created. In this model all the components are deployed in the same CPU. Even though this does not provides information in order to perform partitioning decisions, it gives the possibility to make sure that the model is working correctly

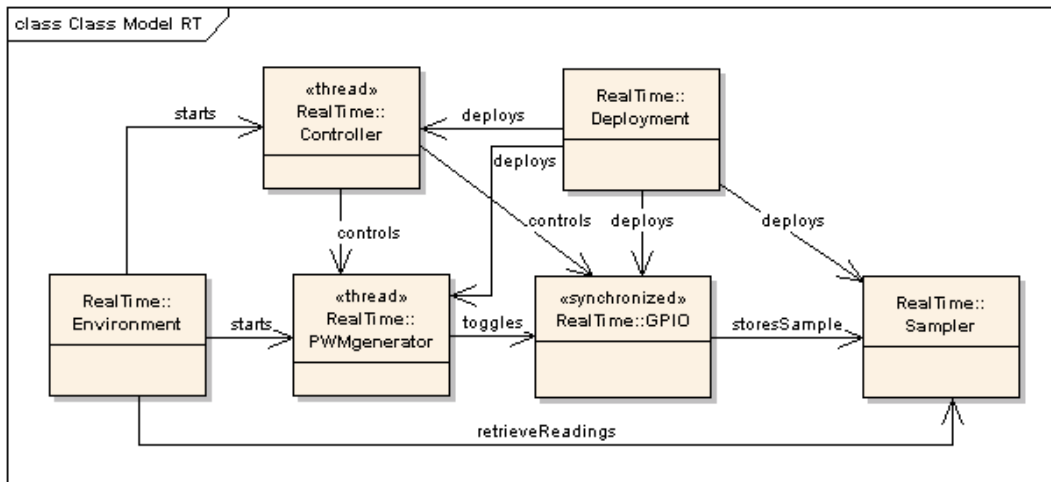


Figure 5.6: Distributed Real-Time model simplified UML class diagram.

under the Overture real-time kernel before moving to the architectural exploration phase (guideline 4).

During the exploration phase, three different architectures have been studied. Two of them are pure software solutions and a third one is making use of a dedicated hardware partition for the PWM signal generation.

5.2.3.1 Architecture 1: an OS-less software based solution

This architecture is purely based on software. The PWM signal generation logic is deployed together with a controller class in a CPU. The GPIO block has been deployed in a separated hardware block communicated at a very high speed (guideline 6) with the CPU software controller. The purpose of such a high bandwidth in the bus used is to model the fast access to a hardware component mounted in the same chip as the CPU. The separated hardware block representing the hardware partition is running at a frequency orders of magnitude above the general purpose CPU frequency (guideline 5).

An overview of the architecture as generated by Overture is shown in figure 5.7.




Figure 5.7: Architecture of the software based solutions.

The deployment of the entities has been carried out as described in the UML deployment diagram shown in figure 5.8.

In order to study the load response of the CPU running the PWM generation logic, the function `timeConsumer` has been introduced. This function is taking processing time progressively each time the controller thread is executed. Therefore, the longer the simulation has been running the more time the `timeConsumer` function will be taking.

Architecture 1: an OS-less software based solution

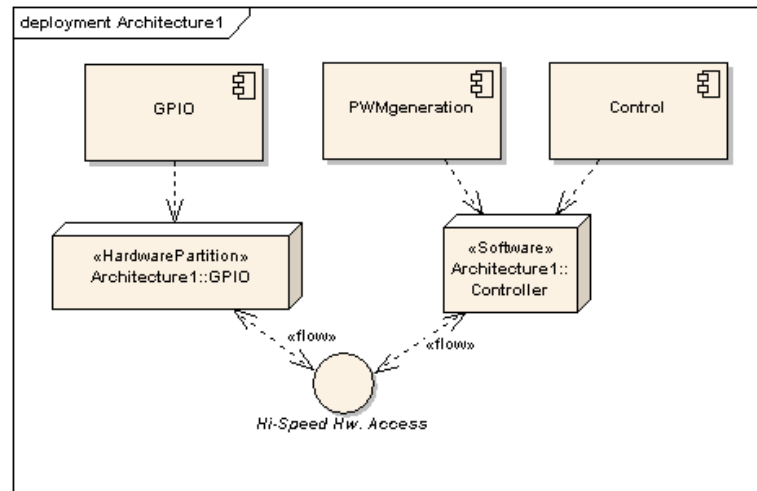


Figure 5.8: Deployment diagram of the software based solutions.

In order to maintain the proper time limits in the signal, duration statements have been used. In this version of the model the `generateCycle` operation in the `PWMgenerator` class has been modelled as shown in the listing below. In this concrete case the target PWM signal has a pulse duration of 2 milliseconds and a low state duration of 18 milliseconds. The `toggleBit` operations are not considered in the time analysis, since the time it takes to toggle the hardware output is orders of magnitude below the phenomena we want to study (generation of a signal in the millisecond range). The increment of the `runCount` has not been considered either, since this is a model dependant parameter that will not be present in the actual implementation.

```

1 duration (2E6) outputInt.toggleBit(1);
2 duration (18E6) outputInt.toggleBit(1);
3 duration (0) runCount := runCount +1;

```

The `PWMgenerator` thread is modelled as a procedural thread that runs until completion. The controller thread is modelled in the same way. The purpose behind this modelling decision is to avoid using periodic constructs, that will be used under an RTOS based implementation.

The generated signal by the OS-less based software solution can be seen in figure 5.9. In this diagram three control pulses can be seen. The duration of the pulses is correctly generated in all three cases. The separation between control pulses is incorrect. This means that the system is failing to generate the control pulses at the required frequency. Note that the time between the second and the third control pulses is considerably higher than time between the first and second pulse. The separation between the pulses keep augmenting while the model execution progresses, making the situation even worse. This means that a pure software based solution would be correct in the case the CPU is not loaded excessively. Load boundaries would depend on concrete platforms, however, the same situation would appear at certain point of load if this control scheme is used. Remark as well that this architecture is the fastest to implement and the cheapest in terms of used silicon area (assuming that a processing unit is already present in the system). Generation of the PWM signal from a OS-less based software solution, while plausible, should be carefully considered due to its load sensitivity.

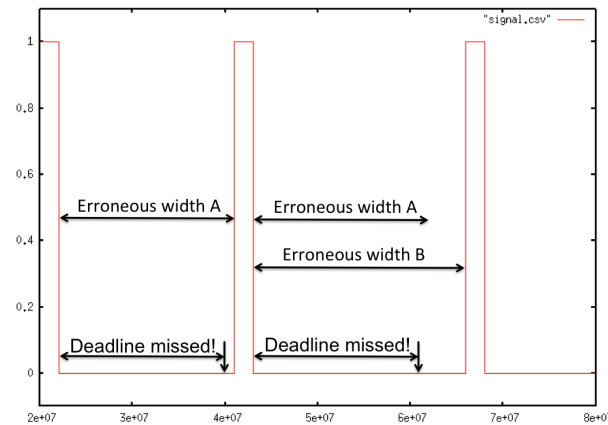


Figure 5.9: Generated control signal by the OS-less software only solution.

5.2.3.2 Architecture 2: a RTOS based software solution

The second architecture evaluated for this case study is making use of a Real Time Operating System (RTOS). The purpose is to exploit the scheduling possibilities that such a software solution offer in order to respect the hard real time deadlines. The used deployment and hardware architecture are exactly the same as the one introduced in section 5.2.3.1. However the scheduling policy in the CPU unit has been altered from *First Come First Served* to *Fixed Priority*. Additionally the time dependant functionality (generateSignal) has been given a higher priority than the controller operations, which are not time critical. Both aspects are illustrated in the following model listing:

```

1 controller : CPU := new CPU(<FP>, 1E9);
2 ...
3 controller.setPriority(PWMunit`generateSignal,10);
4 controller.setPriority(Controller`timeConsumer,1);

```

As it can be imagined, this second software based approach is able to generated the PWM signal properly and respecting the deadlines (see chart on the right side of figure 5.10). It is interesting to study the impact of turning off the priorities and evaluate whether this approach is still valid. In the case no higher priority is given to the generate signal operation, the model is indicating that some deadlines will be eventually missed (see chart on the left side of figure 5.10). Finally, a more detailed view on the correctly generated control signal is depicted on the left side of figure 5.11.

This architecture is more complex that the OS-less software based solution. It is possible to use an open-source RTOS and achieve the expected results. Economically speaking this would not be necessarily more expensive. From the resource consumption point of view, it is clear that this solution requires additional hardware infrastructure (more memory and a dedicated timer to be used by the RTOS). In principle, no additional silicon are would be required.

5.2.3.3 Architecture 3: a combined hardware/software solution

The last evaluated architecture is considering a hardware partition for the generation of the PWM signal, which is show in figure 5.12.

This architecture has been defined in VDM as shown in the next model listing. Since the PWM generation block is deployed in a hardware block, the chosen frequency is orders of magnitude

Architecture 3: a combined hardware/software solution

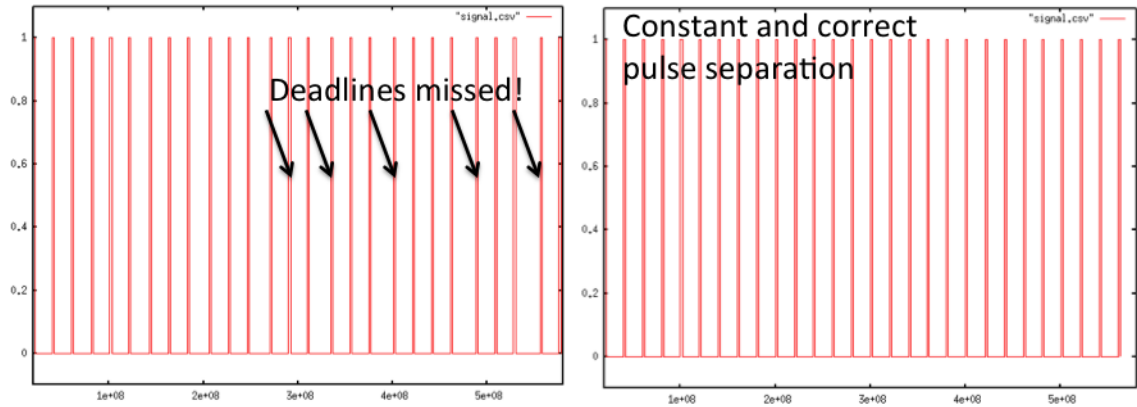


Figure 5.10: Generated signals by the RTOS based software only solution.

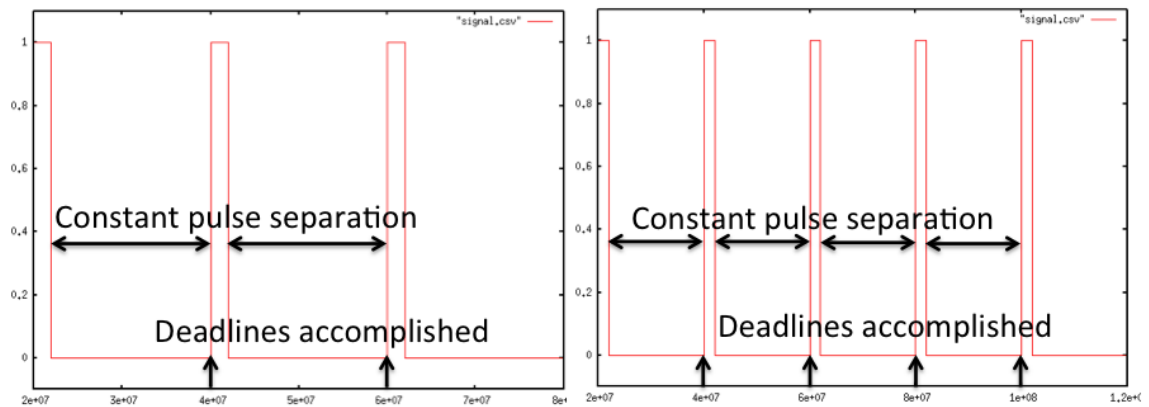


Figure 5.11: Generated signals by the RTOS based with priorities and the Hardware based solutions.




Figure 5.12: Architecture of the hardware based solution.

above the controller frequency (guideline 5). The hardware partition is running at a frequency of $1E9$, while the controller is running at $1E5$. The communication between the CPU and the hardware blocks is established through hi-speed buses, modelling the available high speed interfaces between hardware and software (guideline 6).

```

1 | PWMgenerator      : CPU := new CPU (<FCFS>, 1E9);
2 | controller       : CPU := new CPU (<FCFS>, 1E5);
3 | gpioBlock       : CPU := new CPU (<FCFS>, 1E9);
4 | controlRegister  : BUS := new BUS (<CSMACD>,

```

```

5         72E13,
6         {PWMgenerator,
7           controller}
8       );
9 controlRegister2 : BUS := new BUS (<CSMACD>,
10        72E13,
11        {PWMgenerator,
12         gpioBlock}
13       );

```

The deployment used in this case is depicted in the UML diagram show in figure 5.13. As it can be seen, each entity has been mapped to a separate processing unit.

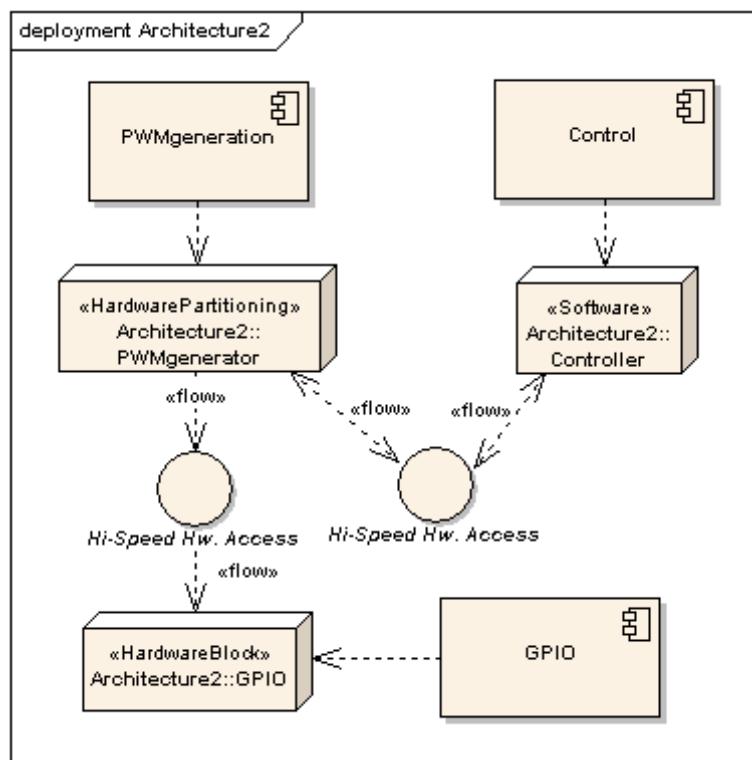


Figure 5.13: Deployment diagram of the hardware based solution.

The logic used is the same as the one explained in section 5.2.3.1. The generated signal is depicted in figure 5.11 on the right. As it can be seen, this architecture is able to generate the PWM signal correctly and within the established time frames. The use of a separate block for the PWM signal generation makes the generated signal independent from the CPU load. The signal is generated correctly no matter how much the CPU is loaded.

This solution is the most adequate regarding performance. On the other hand it requires more silicon area on the FPGA since an additional hardware block has to be deployed on it. In the case an IP block is ready to integration additional development effort would not be necessary. In the case no IP block is available, the required time for VHDL development would have to be considered in the project time budget.

5.3. Implementation of the servo controller

The studied implementations were originally presented in [Speedway11SDK]. The presented results in this section relates the extracted information from the implementation and analyses the results. In addition these results are related to the formal models presented above. The study of the information provided by the actual implementation allows the evaluation of the Hardware/-Software partitioning decisions, and bridges the gap between abstract representation and real implementation (application of guideline 13).

5.3.1 Preliminary details

The models presented in the previous section are making use of a `CPUloader` function that delays the execution at a constant rate of 1 millisecond per iteration. This function is not performing any specific calculation. In this implementation, the `CPUloader` software component has been substituted by an algorithm finding prime numbers. This computational complexity of the algorithm causes a variable execution time per execution. The fact that the `CPUloader` is different does not affect the study of the real-time deadlines fulfilment. The purpose of the models were to illustrate how a delay was affecting the servo control problem, not to model how that delay was created (abstracted detail).

5.3.2 Implementing the conflictive architecture

This architecture is based on a pure software solution. The control logic is running together with the `CPUloader` function on the same CPU and without any kind of operating system. The servo is controlled by the software implementation of the logic described in the sequential VDM model. The function responsible for the `generateCycle` functionality described in VDM is `adjust_servo_manually`. The most relevant parts of this function are the sleep routines `millisleep` and `usleep`, which obtain the appropriate point of time in which the GPIO has to be toggled in order to generate the PWM signal. The variable `high_time_us` is used in order to calculate the duration of the pulse in order to reach the target position, determined by the coded expression in the function. Note that this expression vary depending on the used servo.

```

1 void adjust_servo_manually(int baseaddr,
2                           int servo_number,
3                           int position)
4 {
5     int high_time_us;
6     high_time_us = 700 + ((2000/256) * position);
7     millisleep(25);
8     rc_servo_assert_manual_output(baseaddr, servo_number);
9     usleep(high_time_us);
10    rc_servo_deassert_manual_output(baseaddr, servo_number);
11 }

```

In the code listing presented below, the main control loop is shown. In this loop the performed operations are the servo control operation and the auxiliary cpu loader operation (`findPrime`).

```

1 while (latest_prime_found < highest_prime_to_find)

```

```

2 {
3   adjust_servo_manually(XPAR_AXI_RC_SERVO_CONTROLLER_0_BASEADDR,
4                       1,
5                       0x00 + servo_value);
6
7   latest_prime_found = findPrime(latest_prime_found,
8                                 highest_prime_to_find);
9 }

```

5.3.3 Implementing the optimal architecture

The optimal architecture is offloading the PWM signal generation by the use of the IP block `axi_rc_servo_controller`. The prime number finder is still run in the system processor. The structure of the used IP block can be seen in figure 5.14. The figure shows that the device is clocked, represented by `>0` on the upper left corner, that it has one signal output, represented by `1-` in the upper right corner, and that it is attached to the AXI interface of the microblaze processor. Figure 5.15 shows the device deployed together with other slaves in the final hardware design. As example of other hardware blocks, two GPIO blocks 4-bit with connected to LEDs and switches and an Serial Peripheral Interface Flash memory controller are shown. These blocks are irrelevant to the problem under study, however they give an idea of the variety of devices that can be integrated in the system.




Figure 5.14: IP block implementing the servo control functionality.

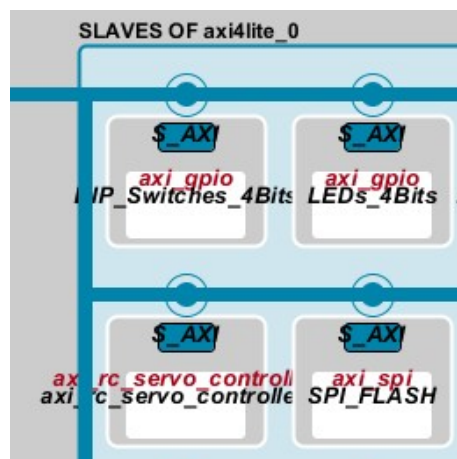


Figure 5.15: IP block deployed.

Even though the use of an IP block abstracts the system engineer from VHDL details, it is worth to inspect how the PWM generation logic is implemented in hardware. The VHDL implementation of the PWM block is shown in the next code listing. The PWM block is organized around a state

Implementing the optimal architecture

machine, which is implemented by the use of a case statement. This state machine presents three possible states: "initial or reset state", "low period" and "high period". A counter is used in order to determine which state should be the next one. The counter is controlled by a clock dependant process. The values the counter may present vary from 0 to `servo_PWM_clock_periods`. This second variable represents the number of elapsed clock periods per PWM cycle. As it can be imagined the number of clock cycles has to be *much higher* than the desired target frequency for the PWM signal in order to get a good resolution. The value the counter is providing is compared with the number of required `low_pulse_width_clock_periods` and `high_pulse_clock_periods` in order to trigger a transition in the state machine or remain in the same state.

```
1 servo_control_out <= '0';
2   reset_control_counter <= '0';
3   case current_state is
4   when reset => reset_control_counter <= '1';
5     next_state <= low_period;
6
7   when low_period =>
8     if (control_counter >= low_pulse_width_clock_periods) then
9       reset_control_counter <= '1';
10      next_state <= high_period;
11     else
12       next_state <= low_period;
13     end if;
14
15   when high_period => servo_control_out <= '1';
16     if (control_counter >= high_pulse_width_clock_periods) then
17       reset_control_counter <= '1';
18       next_state <= low_period;
19     else
20       next_state <= high_period;
21     end if;
22 end case;
```

```
1 control_counter_process : process (clk, rst)
2 begin
3   if clk'event and clk = '1' then
4     if reset_control_counter = '1' then
5       control_counter <= 0;
6     elsif (control_counter = servo_PWM_clock_periods) then
7       control_counter <= 0;
8     else
9       control_counter <= control_counter + 1;
10    end if;
11  end if;
12 end process;
```

The most interesting conclusion that can be drawn from this analysis, is that *the logic behind the PWM generation in VHDL is the same as the one modelled in VDM*. Besides the control logic, the VHDL implementation contains the required logic to interact with the processor. The IP servo control block is integrated as a slave hardware device, mapped to a control register in the microblaze processor. From the software side, the servo is controlled by invoking the function `rc_set_servo_position`, which will take as a parameter the register address, the servo number and the target position. The IP hardware block will be continuously polling the associated register and repositioning the servo according to the hold value. This process is run in parallel and independently of the state of the software functionality running in the microblaze CPU.

```

1 void rc_servo_set_servo_position(int baseaddr, int servo_number,
2 int position)
3 {
4     servo_number--;
5     if (servo_number >=0)
6         rc_servo_set_servo_register(baseaddr,
7                                     8 + (servo_number*4),
8                                     position);
9 }

```

5.3.4 RTOS based solution

A third possible implementation, which has been explored from the VDM-RT context, is the integration of a Real Time Operating System. This third design alternative has been studied but not implemented in the FPGA. However, similar results to the ones exhibited by the VDM-RT model can be expected. The considered RTOS for this implementation have been *Xilikernel* and *FreeRTOS*. *Xilikernel* is not a viable option for new designs or research activities. Xilinx representatives have stated in informal conversations that *Xilikernel* is not a priority for Xilinx when it comes to development effort and maintenance. As an alternative they highly recommend the use of *FreeRTOS*³.

As it was explored in the RTOS based VDM model, the generation of the control signal could be incorporated in a periodic thread that is executed at a period of 20ms. A thread (called task in *FreeRTOS*) can be used as shown in the next code listing. One of the most relevant parameters of the thread creator is the function entry point. In this case, the entry function will be associated to the generate signal function, as it was shown in the models. The execution of this task will be started by a software timer, which will be generating an interrupt every 20ms. The task should be executing the `generateSignal` functionality within a *critical section*, so a preemptive context switch will not occur. Interrupts should be deactivated *before* entering the critical section.

```

1 portBASE_TYPE xTaskCreate(
2     pdTASK_CODE entryPoint,
3     const portCHAR * name,
4     unsigned portSHORT usStackDepth,
5     void * Parameters,
6     unsigned portBASE_TYPE uxPriority,

```

³FreeRTOS is open-source and is available for download at <http://www.freertos.org/>.

Evaluating the implemented architectures

```
7 | xTaskHandle * callback // If needed
8 | );
```

The next code listing shows the thread signature that should be used in a FreeRTOS based implementation.

```
1 | xTaskCreate (
2 |     generateSignal,
3 |     "RCservoControl",
4 |     SMALL_SIZE,
5 |     (void*)targetTime,
6 |     HIGH_PRIORITY,
7 |     NULL
8 | );
```

While possible, this software implementation can be highly inefficient, since the system can be busy generating the PWM signal up to 10% of the time. Besides performance considerations, a more critical situation could be that other real-time events are not being processed because the PWM task is being served.

5.4. Evaluating the implemented architectures

In this section two implementations will be evaluated. These architectures correspond to the OS-less software solution and to the combined hardware/software solution. The results provided by this evaluation were initially predicted by the models presented above, which were used to make the partitioning decisions of the system. In a real development setup, only the most suited architectural candidate will be selected for implementation. In order to validate the thesis that VDM-RT models can provide the necessary information to make partitioning decisions, the two opposite architectures have been evaluated. As it will be presented below, the performance of the different solutions was properly anticipated by the abstract VDM-RT models.

5.4.1 Evaluation of the conflictive architecture

Xilinx Software Development Kit provides a tool to profile the execution of a certain application. The advantage of the Xilinx profiler is that the evaluation is performed when the application is running on the actual hardware, providing real measurements instead of predictions. The profiler has been used in order to evaluate the performance of the "software only" implementation. In order to get more meaningful results, the serial traffic has been eliminated by commenting out the print statements in the application. The reason behind this is that a considerable amount of time is used in order to send information through the console. This process may mask the potential bottlenecks that may appear in the real application. After eliminating the print statements and some preliminary profiler setups, a test run can be executed. This test run offered the results shown in figure 5.16.

The time critical functions `adjust_servo_manually` and `calculate_servo_position` are taking 0.26% and 0.49% of execution time respectively. On the other hand the CPU loading function `findPrime` is requiring 9.75% of execution time, and invoking arithmetic functions that

Name (location)	Samples	Calls	Time/Call	%Time
Summary	1097198			100,0%
??	7			0,0%
_profile_clean.c	0			0,0%
_profile_init.c	1			0,0%
_profile_timer_hw.c	3			0,0%
complete_application.c	120580			10,99%
adjust_servo_manually	2834	100	283.397us	0,26%
calculate_servo_position	5346	100	534.594us	0,49%
findPrime	106978	99	10.805ms	9,75%
main	4135	0		0,38%
my_interrupt_handler	75	340	2.205us	0,01%
printPrime	1212	99	122.423us	0,11%
crt0.S	13			0,0%
libgcc2.c	272848			24,87%
profile_cg.c	0			0,0%
profile_hist.c	149			0,01%
sleep.c	178806			16,3%
udivsi3.asm	257471			23,47%
umodsi3.asm	264184			24,08%
xgpi.o	3066			0,28%

Figure 5.16: Software based solution profiling.

are requiring almost 50% of processing time if considered together (hudi`divsi3` and `umodsi3`). This shows that the software bottleneck is introduced by the CPU loading function `findPrime`. This function is making the control function `adjust_servo_manually` missing its real time deadlines. The profiler shows how the problem predicted by the VDM-RT models is present in the actual system implementation.

Besides detecting the conflictive situation by the use of the profiler, it is possible to monitor the generated PWM signal with a scope, and study how it changes over time.

During the first seconds of execution, it can be seen how the PWM signal is generated properly, respecting the time limits for the pulse and at the required frequency. Figure 5.17 shows both constraints on the left and right scope dumps respectively.




Figure 5.17: High pulse duration and correct signal period.

After the system has been running for a period of time of approximately 20 seconds, it is possible to appreciate the problems that arise if the software only architecture is used. Figure 5.18 on the left shows how the spacing between two consecutive control pulses is taking up to 1.3 seconds, orders of magnitude greater than the initial 25 ms deadline. Moreover, since the CPU loader is consuming processing time at a variable rate, the delay introduced in the control pulse generation

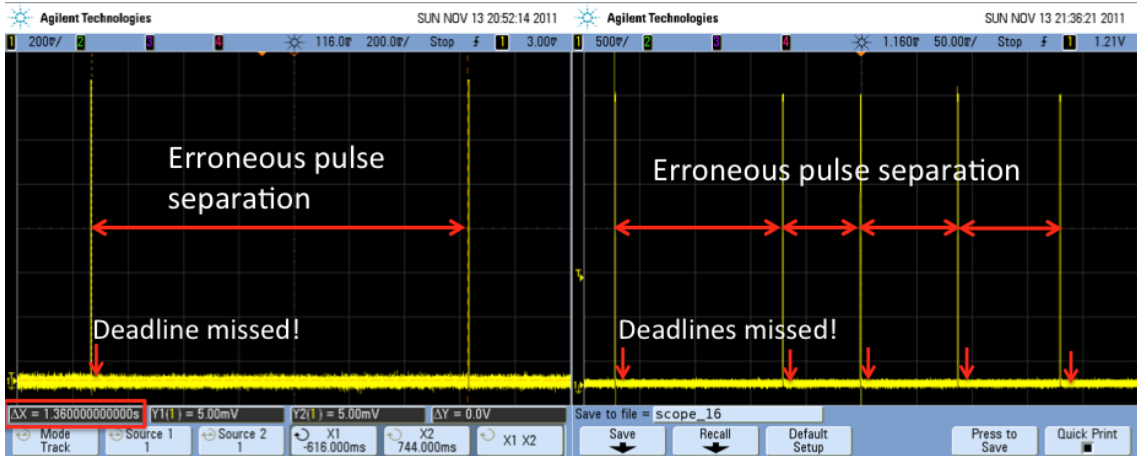


Figure 5.18: Generated control signals after 20 seconds of operation.

is varying as well. This makes the control error variable during system operation. Uneven spacing between the control pulses can be seen in the chart on the right in figure 5.18.

5.4.2 Evaluation of the optimal architecture

The optimal architecture, which makes use of a hardware block for the PWM signal generation, has not been profiled. This second solution has been directly deployed into the FPGA and its generated signal monitored with the scope. An overview of the generated signal can be seen in figure 5.20. As it is depicted, the spacing between the control pulses seems to be correct and evenly distributed. The measurement was repeated at different moments of time over several minutes with no changes.

In order to take a closer look at the signal, the timebase was decreased. This allowed to analyse the duration of the control pulse and the exact separation between consecutive pulses. The duration of the control pulse, shown in figure 5.19 on the left is correct, presenting a value of 1.76ms. The separation between two consecutive control pulses, shown in figure 5.19 is correct as well, presenting a value of 29.4 ms.

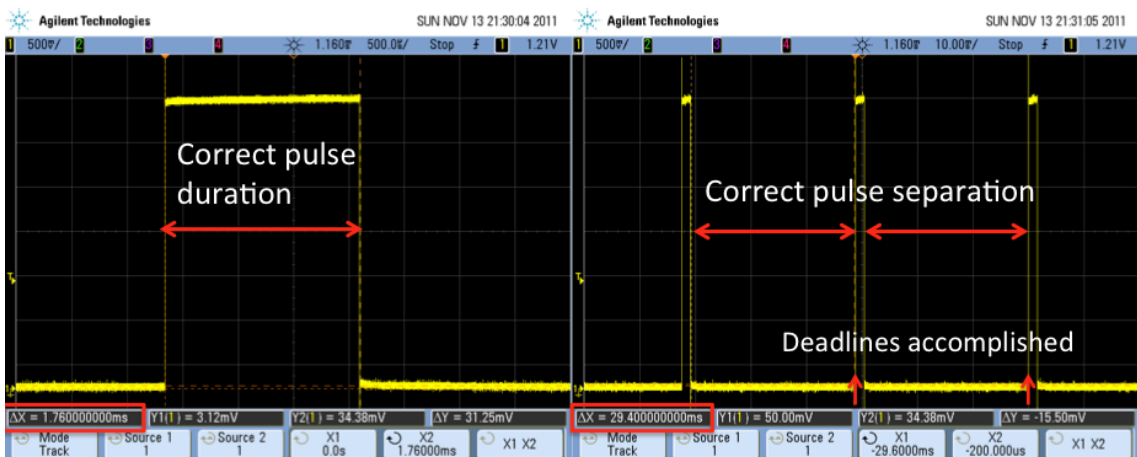


Figure 5.19: Detailed view of the generated PWM signal by the IP PWM block.




Figure 5.20: Overview of the generated PWM signal by the IP PWM block.

5.4.3 Physical results

In order to see the results of both architectures physically, a servo was connected to the GPIO pin in the FPGA. When the software only architecture was in use, it was possible to see the jitter in the servo axis after a short period of operation. The jitter was increasing considerably until a full oscillation between 0 and 180 degrees was reached. Remark that a continuous operation under this control signal can damage the servo itself. If the servo is performing a critical control, the consequences of missing the real time deadlines can be serious depending on the concrete application. The hardware based architecture was physically studied as well. As it was predicted by the VDM RT models and shown by the scope analysis, the target position is maintained constant making impossible to move the servo axis. In this case, a complete implementation was possible to be carried out due to the existing literature, material and the simplicity of the case study. However, there are different control cases that are challenging, where different logical control strategies and partitionings might have an impact on the physical performance of the solutions. These more complex cases can be harder to implement, and a model-based setting is the only way to perform the design space exploration required in order to perform a sound partitioning. Evaluation of partitionings and performance of different control logic can be performed by VDM (discrete event modelling). On the other hand, the study of a continuous time process is hard to perform under a modelling environment like VDM. In this kind of scenario, a tool like DESTTECS can be a good solution in order to contextualize partitioning decisions and analyse their impact on the physical environment. Additional information on the DESTTECS EU FP7 project can be found on [Broenink&10].

5.5. Summary

This chapter has presented how a VDM-RT based modelling methodology can be applied in the hardware software co-design of embedded systems. The methodology has helped to evaluate several architectures and decide on a concrete implementation, making use of IP-cores. Finally, the system implementation has been carried out. Feedback from that concrete implementation has been obtained through application profiling. With demonstration purposes, a conflictive architec-

Summary

ture has been implemented as well. Correlation between created models and actual results has been shown.

The VDM models have helped on creating an unambiguous representation of the system, that can be executed under different scenarios. By applying modelling it has been possible to represent the hardware/software co-design partitioning issues and analyse the performance regarding real-time deadline accomplishment. Even though the servo case is simple from a technical point of view, creating models that provide the same kind of information as the one obtained from VDM models using, for example, SysML, could have been more complex.

Case study: Development of an AVB endpoint device

This chapter is focused on a complex industrial case study, the hardware/software co-design of an Audio Video Bridging (AVB) endpoint device. The methodology proposed in chapter 4 will be applied. This chapter is structured in seven sections. An introduction is presented in section 6.1. An overall description of the AVB protocol will be given in section 6.2. The relevant aspects for the hardware partitioning of the system will be presented in section 6.3. The created models for the hardware/software co-design of the system will be presented in section 6.4. Evaluation of different architectures will be described in section 6.5. The outcome of the model would be briefly evaluated in section 6.6. Finally, section 6.7 summarizes the chapter.

6.1. Introduction

This chapter presents a second case study proposed by the home electronics manufacturer Bang & Olufsen. The focus of this case is the specification and analysis of an Audio Video Bridging compliant endpoint device. The preliminary analysis will create the basis for the hardware software co-design of the system, and it will assist in making the initial partitioning decisions of the system.

6.2. Description of the AVB protocol

The AVB protocol is used in the transmission of multimedia content over AVB compliant networks. It provides the necessary mechanisms to ensure the **real-time deadlines** associated with this kind of content, as well as the **bandwidth reservation** and the facilities to avoid **traffic bursting**. The AVB protocol is a layer 2 transport protocol according to the Open Systems Interconnection (OSI) classification. This protocol is specified in the standard [IEEE-std1722-2011]. This standard specifies that in order to create an AVB compliant device the following protocols have to be supported:

802.1AS: time synchronization protocol. This protocol is responsible for distributing a common notion of time among all the nodes that belong to the AVB network.

802.1Qav: stream reservation protocol. This protocol is responsible for allocating 75% of the available network bandwidth for transmission of AVB traffic.

802.1Qat: traffic shaping protocol. This protocol ensures that the traffic is sent over the network smoothly and at a constant pace.

IEEE1722 can be used on top of multiple media layers. These physical mediums include optical fibre, wireless networks and Ethernet. The main development effort at the time of writing is focused on the 802.3 physical layer (Ethernet). Note that AVB has nothing to do with TCP/IP. TCP/IP traffic can be used in parallel to IEEE1722, for example, in order to send control commands. TCP/IP is not necessary in order to stream multimedia content with AVB. A contextualized view of the involved protocols and layers is depicted in figure 6.1. An overview of AVB and its relation to these protocols is given in [Garner&07].




Figure 6.1: Protocol stack showing the involved protocols in AVB. Source: [IEEE-std1722-2011]

Due to its timing performance, the AVB protocol application is not limited to the transmission of multimedia content. Current developments include industrial automation and robotics. Even though these areas are currently being researched, the immediate business case for AVB is the home entertainment market.

6.2.1 The traffic shaping protocol 802.1Qav

The traffic shaping protocol is known as well as a "queuing and forwarding protocol". All the network equipment in an AVB network must be compliant with the traffic shaping protocol. This protocol is specified in [IEEE-Std802.1Qav-2009]. The purpose of this protocol is to provide a uniform traffic at a constant pace. By applying traffic shaping bursting traffic is avoided. A second consequence of its application is that the receiver buffers are not overflowed. The traffic shaper creates several time slots and assigned them to AVB and other traffic (for example, TCP/IP traffic) depending on the listener capabilities and the time deadlines. This scheme resembles the Time Division Multiplexing (TDM) technique used in many communication systems. Figure 6.2 shows the basic structure of a traffic shaping unit. The shaper is composed by two multiplexers, a Priority-Based multiplexer and a Timing aware scheduler. The latter is the only multiplexer with access to the physical medium. The Priority-based scheduler manages non-critical data, such as TCP/IP traffic. This traffic is delivered as best-effort to the Timing aware scheduler. The queues for critical traffic, for example AVB traffic, are labelled as *Class 4/5 queues*, and are directly connected to the final multiplexer. Finally, the timing aware scheduler classifies the traffic and is responsible for putting the data on to the physical medium when it is scheduled. Further details on the protocol operation are out of the scope of this thesis work. For additional information refer to [Garner&07] and the standard [IEEE-Std802.1Qav-2009].

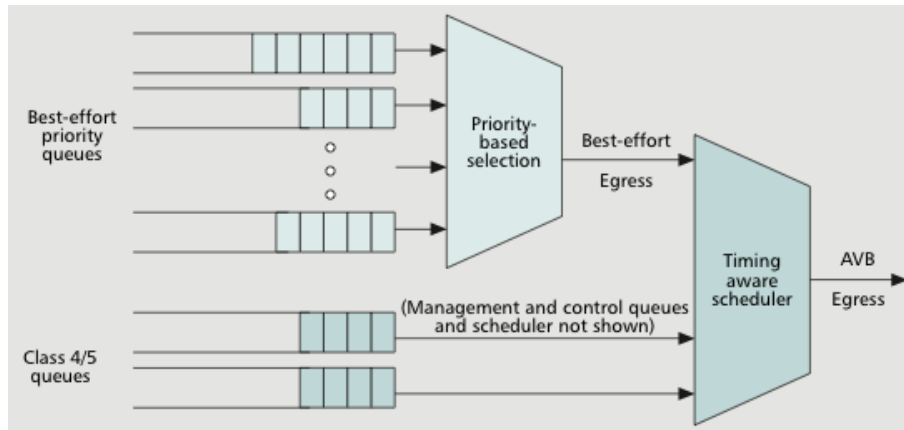


Figure 6.2: AVB and TCP/IP traffic managed by packet filtering. [Garner&07]

6.2.2 The stream reservation protocol 802.1Qat

AVB is effective in time-sensitive content transmission in part due to the network resources that are allocated for its operations. This process is done by applying the stream reservation protocol, specified in [IEEE-Std802.1Qat-2010]. The application of the stream reservation protocol makes it possible to minimize jitter and guarantee a *deterministic* low latency. All the network equipment in an AVB network must be compliant with the stream reservation protocol in order to enable operation. The stream reservation protocol uses the Virtual Local Area Network (VLAN) tagging mechanism in order to filter the network packets before they are transferred to the physical medium. This implies that the packet filtering process is carried out in a very low layer of the communications stack, and can be done by using bitwise operators. The process of configuring the network trees is more complex and involves a dialogue between the talker (information source) and the listener (information sink). Additional details on protocol setup, maintenance and operation are out of the scope of this work. For further details refer to [Garner&07] and the standard [IEEE-Std802.1Qat-2010].

6.2.3 The time synchronization protocol 802.1As

This protocol is specified in [IEEE-Std802.1As-2011]. According to this standard, a device that is able to receive or generate time corrections is called **Time Aware System (TAS)**. Each TAS is considered as a system by itself. Therefore, the term system should not be associated with the set of devices conforming the network. A network composed by several TAS could be considered as a network of systems.

6.2.3.1 TAS Structure

Each TAS is composed by several modules. These modules are depicted in figure 6.3. A brief description of each module is presented in the following lines:

Media Dependant Layer: is responsible for the access to the physical medium. Functions that vary depending on the used physical connection are implemented at this layer.

Port Sync: is the information entry point for the TAS. Each Port Sync is attached to a Media Dependant Layer access point.

Site Sync: retrieves the information from every Port Sync structure present on the system. It coordinates the time information dissemination depending on the port functionality. It acts as an interface between the registered ports and the Clock Slave and the Clock Master.

Clock Master: is present only in the case a TAS is Grand Master capable. In the case the TAS is the Grand Master, the Clock Master will be generating the time information that will be sent to the rest of the TAS present in the network. In the case the TAS is not acting as Grand Master, the Clock Master will be registering the difference between the time information received and the time information generated by itself. In the case the current Grand Master leaves the network, and this TAS assumes that role, then it is aware of the existing difference in time, avoiding potential discontinuities.

Clock Slave: is responsible for receiving the time information from the Site Sync entity and processing it. The time information will be made available to the time consuming application through the interfaces implemented by the Clock Slave. The time information will be notified to the Clock Master in the case the TAS is not acting as the network Grand Master, in this case the above explained offset will be computed by the Clock Master.

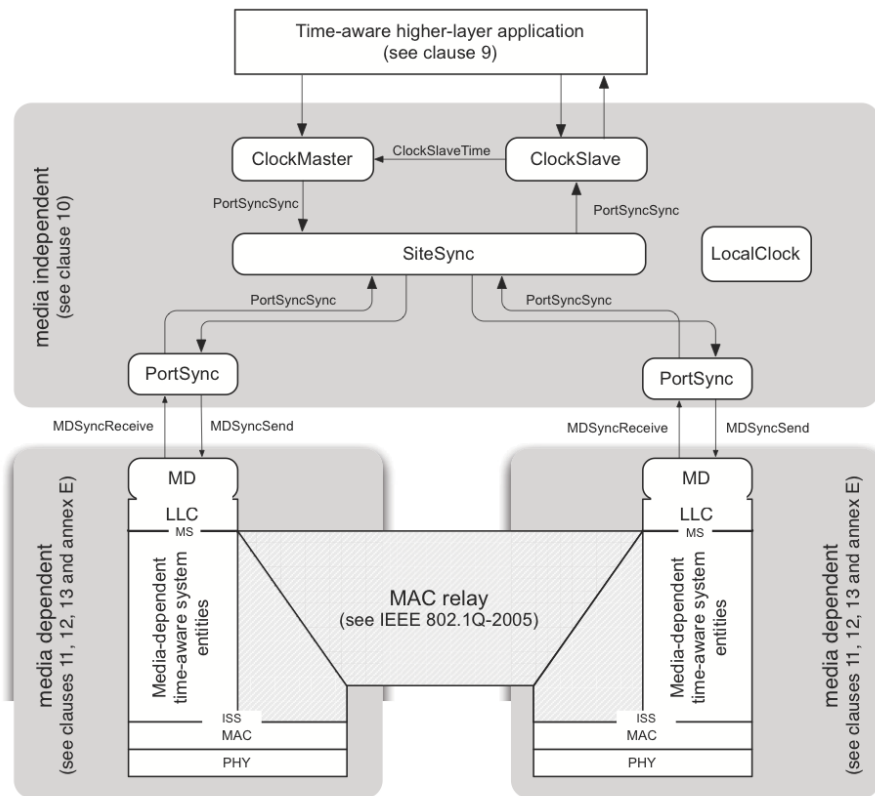


Figure 6.3: Overall structure of a Time Aware System. [IEEE-Std802.1As-2011]

Ports may present different roles: *Master*, *Slave*, *Passive* or *Disabled*. The Master Ports in a TAS are sending the information to other TAS, which are listening in the Slave Ports. Therefore, there is a one-to-one association between each Master and Slave port. Passive and Disabled ports are used to avoid loops in the network topology. These kind of ports are not playing an active role in the network operation. Figure 6.4 shows a Grand Master Time-Aware Bridge connected to a

TAS Behaviour

Time-Aware Bridge. As it is shown all the ports of the Grand Master are acting as Master ports. The second port is connected to the slave port of the Time-Aware Bridge. This second device is forwarding time information over the two lower master ports. Left and right ports are configured as Passive ports, and are not sending nor receiving time information.




Figure 6.4: A Time-Aware Bridge connected to a Grand Master. [IEEE-Std802.1As-2011]

6.2.3.2 TAS Behaviour

The time synchronization is a two step process in which two separated messages have to be sent. These messages are the *SyncEvent* and the *FollowUpMessage*.

SyncEvent: signals the point of time in which the time information is sent to a slave.

FollowUpMessage: contains the necessary information to compute the point of time in which the SyncEvent was originated at the Master. The follow up has to be generated after the Sync Event was sent, so the precise timestamp of the sending event can be obtained. Further details will be provided in the example below.

The FollowUpMessage is composed by the following fields:

Precise Origin Timestamp: is the original timestamp generated at the Grand Master when a SyncEvent was sent to the Media Dependand Layer.

Cumulative Rate Ratio: is the ratio between the frequency of the master TAS and the slave TAS. Each TAS will generate a different cumulative rate ratio. Further details will be presented in the coming example.

Correction Field: time correction computed at every node after forwarding the SyncEvent. Further details will be presented in the coming example.

Remark that the Precise Origin Timestamp remains constant, and it is never altered by the TASs that are sending a modified Follow Up Message. Only the information contained in the Cumulative Rate Ratio and Correction Field is altered. An example is given below in order to give a better explanation of the synchronization process.

The example network is composed by four TASs. The first system is acting as Grand Master. The remaining three systems are connected one after another (daisy chained). Device two and three are *bridges*, since they can communicate with more than one TAS and forward information. The

devices are deliberately operating at different frequencies to make the example more illustrative. A deployment diagram of this network is shown in figure 6.5. Frequency is specified the generic unit pulses per time unit (ppt). Time is specified in time units (tu).




Figure 6.5: Example network scenario.

The synchronization process starts with the slave ports measuring the propagation delay and retrieving the frequency of the devices with master ports. Once the frequencies for each master device associated to each slave port have been determined, the *Neighbour Rate Ratio* for each TAS is computed. The Neighbour Rate Ratio is defined in equation 6.1.

$$NeighbourRateRatio_{TAS\#n} = \frac{Freq_n}{Freq_{n-1}} \quad (6.1)$$

The product of two or more Neighbour Rate Ratio will produce the Cumulative Rate Ratio. This figure allows to compare time measurements carried out in a certain node in relation to the Grand Master. The Cumulative Rate Ratio is computed by applying the equation 6.2.

$$CRR_{TAS\#n} = \frac{Freq_{TAS\#1}}{Freq_{GM}} \cdot \prod_2^n \frac{Freq_{TAS\#n}}{Freq_{TAS\#n-1}} \quad (6.2)$$

Which is equivalent to equation 6.3.

$$CRR_{TAS\#n} = \prod_1^n NeighbourRateRatio_{TAS\#n} \quad (6.3)$$

Considering TAS # 2, its Cumulative Rate Ratio would be computed as shown in equation 6.4.

$$CRR_{TAS\#2} = \frac{Freq_{TAS\#1}}{Freq_{GM}} \cdot \frac{Freq_{TAS\#2}}{Freq_{TAS\#1}} = \frac{Freq_{TAS\#2}}{Freq_{GM}} \quad (6.4)$$

Any measurement performed at the TAS #2 can be multiplied by the Cumulative Rate Ratio computed for that TAS, having as a result a time measurement related to the time base used in the Grand Master.

Table 6.1 shows the result of applying systematically these expressions to the Time Aware Systems deployed in the example network.

After a delay of 10 time units, introduced by the physical link layer, the Sync Event sent by the GM is received and timestamped by the TAS#1 producing the *IngressTimeStamp*. The Sync Event is

TAS Behaviour

	GM	TAS#1	TAS#2	TAS# 3
Freq	30	10	20	10
Neighbour Rate Ratio	1	10/30	20/10	10/20
Cumulative Rate Ratio	1	$1 * 10/30$	$1 * 10/30 * 20/10$	$1 * 10/30 * 20/10 * 10/20$
	1	$= 1/3 = 0.33$	$= 2/3 = 0.66$	$= 1/3 = 0.33$

Table 6.1: Neighbour Rate Ratio and Cumulative Rate Ratio computation

GM generated Follow Up Message	Values
Precise Origin Timestamp	A
Cumulative Rate Ratio	1
Correction Field	0

Table 6.2: Grand Master Follow Up Message

followed by the Follow Up Message. TAS#1 forwards the Sync Event and timestamps the sending process, generating the *EgressTimeStamp*. By calculating the difference of these two timestamps the *Residence Time* for TAS#1 is determined. The Residence Time indicates the time the system has used on receiving, processing and forwarding the event. Residence time is calculated as follows:

$$ResidenceTime_{TAS\#n} = EgressTimeStamp_{TAS\#n} - IngressTimeStamp_{TAS\#n} \quad (6.5)$$

The ResidenceTime has to be taken into consideration so the hops can be compensated when calculating the original GM time. At this point, it is possible to create the Follow Up message contents that will be sent to TAS#2. As explained above, the Precise Origin Time Stamp will not be altered. the Cumulative Rate Ratio has been previously determined. The Link Delay has been measured by the slave port. Finally, the residence time has been obtained applying the last expression. The expression 6.6 links all this variables together, and enables the calculation of the new Correction Field.

$$CorrectionField_{TAS\#n} = ResidenceTime_{TAS\#n} \cdot CRR_{TAS\#n} + LinkDelay + \quad (6.6) \\ + CorrectionField_{TAS\#n-1}$$

Applied in the case for TAS#1:

$$CorrectionField_{TAS\#1} = 20 \cdot \frac{10}{30} + 10 + 0 = 16.66 \quad (6.7)$$

Finally, the follow up message is sent to the TAS#2. The second TAS will repeat the same procedure as explained above. A special case is presented by the TAS#3. Since this TAS is the limit of the system, an endpoint device, it does not forward the time information. Therefore it does not create a Follow Up Message. TAS#3 consumes the received follow up information and associates this data to the generated timestamp when the sync event was received locally. With the contents of the Follow Up Message, any TAS is able to determine the GM time that corresponds to the moment of time in which the Sync Event is received locally and timestamped. In order to make this example more clear, the Follow Up Message related calculations for all the TAS in the network are shown in tables 6.2, 6.3, 6.4.

TAS#1 generated Follow Up Message	Values
Precise Origin Timestamp	A
Cumulative Rate Ratio	10/30
Correction Field	$20 * 10/30 + 10 + 0 = 16.66$

Table 6.3: TAS#1 generated Follow Up Message

TAS#2 generated Follow Up Message	Values
Precise Origin Timestamp	A
Cumulative Rate Ratio	$10/30 * 20/10 = 20/30$
Correction Field	$5 * 20/30 + 10 + 16.66 = 29.99$

Table 6.4: TAS#2 generated Follow Up Message

Details on the timing performance of the 802.1As can be found in [Garner&09]. For a more thorough description of the protocol refer to [Garner&11]. Do not read [IEEE-Std802.1As-2011] before having an overall understanding of [Garner&11]. The standard document is appropriate in order to find specific technical details on the protocol. On the other hand, the same document fails to give an overall simple explanation.

6.3. Partitioning decisions and modelling the AVB compliant device

In order to model an AVB compliant endpoint device, the following aspects have to be taken into considerations:

6.3.1 Protocol modelling

An AVB compliant device has to provide an implementation of the protocols presented in the previous section: the *traffic shaping protocol*, the *stream reservation protocol* and the *time synchronization protocol*. Both traffic shaping and stream reservation protocols are heavily based on operations over data queues. If the case under study is the SoS and network perspective, it makes more sense to build models by using *queue theory*. Such a modelling approach could provide a mathematically proven architecture with the optimized number of queues, frames per second and nodes in the network. In this sense, it is clear that VDM-RT is not the most promising modelling technology. Taking in consideration a single node in the network, the implementation of the traffic shaping and stream reservation protocols is reduced to a set of packet filtering operations over the hardware buffers. These are mostly bitwise comparisons over the buffers, that will determine when the Ethernet packets will be sent. Since both protocols are amendments to the widespread 802.1Q tagging protocol, they are available and integrated in COTS Ethernet interfaces. From the partitioning perspective it is clear that there are highly efficient hardware implementations in the market ready for system integration.

6.3.2 Real-time characteristics

Considering the real time perspective the time synchronization protocol is more challenging than traffic shaping and stream reservation protocols. The real time analysis possibilities offered by Overture will allow a coarse grained study of the system time performance. The most relevant issues to analyse for a slave endpoint device, considering time synchronization, are:

Introduced error at Sync time stamping: The sync event is time stamped when it is received by the endpoint AVB compliant device. The time difference between the reception of the Sync event and the point of time in which the time stamping takes place is the introduced time error in the device. It is desirable to introduce the least error possible. The error budget for a set of commercial WiFi speakers is set to +/- 10 microseconds [Stanton08].

Sync processing time: The performance of the device in processing sync events determines the maximum rate at which it can be receiving them. This limit should not be overridden, otherwise bandwidth will be wasted¹.

Note that not incorporating traffic shaping and stream reservation protocols in the time synchronization models will not have any impact on the obtained results. Time synchronization related traffic will always be treated as a high priority traffic, with access to the physical channel as soon as it is requested by its implementation. Therefore, it is fair to model the buses as dedicated point to point connections between the connected devices.

6.3.3 Partitioning factors to be considered by the models

The VDM-RT models will help on evaluating the **precision** offered by different architectures. This factor will be measured by the error introduced in the different designs. The lower the introduced time error the more precision the architecture is offering. The aim is to provide executable models that will act as enabling tool to study the **time feasibility of several architectures**. The different architectures will be composed by a variety of hardware and software components. It is worth to keep in mind that, as exposed in chapter 2, a certain partitioning decision will be accompanied by a set of trade-offs in terms of: **used silicon area, development time, flexibility of the design, scalability and price**.

It is the intention that the proposed VDM-RT model will help on finding the best compromise for a solution depending on the target application.

6.4. Modelling the AVB time synchronization protocol

6.4.1 Sequential model

The initial sequential model of the system gives the possibility of identifying the main entities involved in system operation. The main outcome of the sequential model is to represent the structure and the time independent behaviour of the system. This model is incorporating a very limited notion of time, represented as *timesteps*. Even though the notion of time introduced in this model is not providing the mechanisms for timing analysis, it is sufficient in order to show the event ordering. Abstracting time details is beneficial in an initial modelling approach, since it allows the modeller to focus only on designing the system structure and distributing the

¹In a poorly designed system, even overflows may occur.

responsibilities among the components. An additional outcome is that properties that should be exhibited by the system can be incorporating at this stage. Finally, the application of modelling gives the possibility of reaching a good understanding of the protocol. In the case of complex protocols, like the one considered in this section, the right model will be created after several attempts (guideline 2).

The sequential model of the system has served to detect the following entities in the time synchronization subsystem:

App: represents the time consuming application.

ClockSlave: represents the clock dependant on the time corrections received by the TAS. It provides the time information to the App entities.

SiteSync: represents the intermediate layer between the received information in the ports and the ClockSlave entity.

Port: represents an abstract entity that can be acting as a Master or as a Slave port depending on the device configuration.

Slave: represents a particular kind of port, used to received incoming time corrections.

Clock: represents the local system clock. It is used for timestamping and as local time reference for the TAS using it.

PhyMessage: represents the buffer receiving the sync event and the follow up message. This class acts as a proxy between the endpoint and the external devices (application of guideline 10).

Figure 6.6 shows the UML class diagram of the sequential model. The environment class has been omitted in order to make the diagram more readable. The information interface is represented by the `PhyMessage` class, which is associated to the abstract class `Port`. A concrete instance of the `Port` class should be used in order to manage the information exchange. The `Clock` class is associated to the `Ports` so a timestamp can be applied to the incoming messages. The class `Slave` is the concrete implementation used in this model. The `Ports` are managed by the `SiteSync` entity. The received time corrections are provided to the `ClockSlave` attached to the `SiteSync` entity. This clock will be responsible for forwarding the corrections to the applications requiring time information.

In order to reach a better understanding of the protocol, a generic deployment that involves bridges and endpoint devices has been modelled. In this deployment concrete implementations of the `Port` class, in the form of Master and Slave ports have been used. Figure 6.7 shows the object diagram of the created model. Even though the focus of this case study is the development of AVB compliant endpoints (TAS1 and TAS4 in figure 6.7), this deployment models as well the use of AVB compliant bridges (TAS2 and TAS3 in figure 6.7). Such a model does not bring additional information about how an endpoint device should be partitioned. However, it is possible to contextualize the operation of the endpoint device in a more concrete and generic scenario. Only after thoroughly understanding how the target system is interacting with the environment, makes sense to move forward towards the more focused modelling.

6.4.2 Concurrent model

The purpose behind creating a concurrent model of the system is the same as the one exposed in section 5.2.2. It is worth to mention that the notion of time in this model remains unaltered with respect to the sequential model. In order to represent the concurrent progression of time the

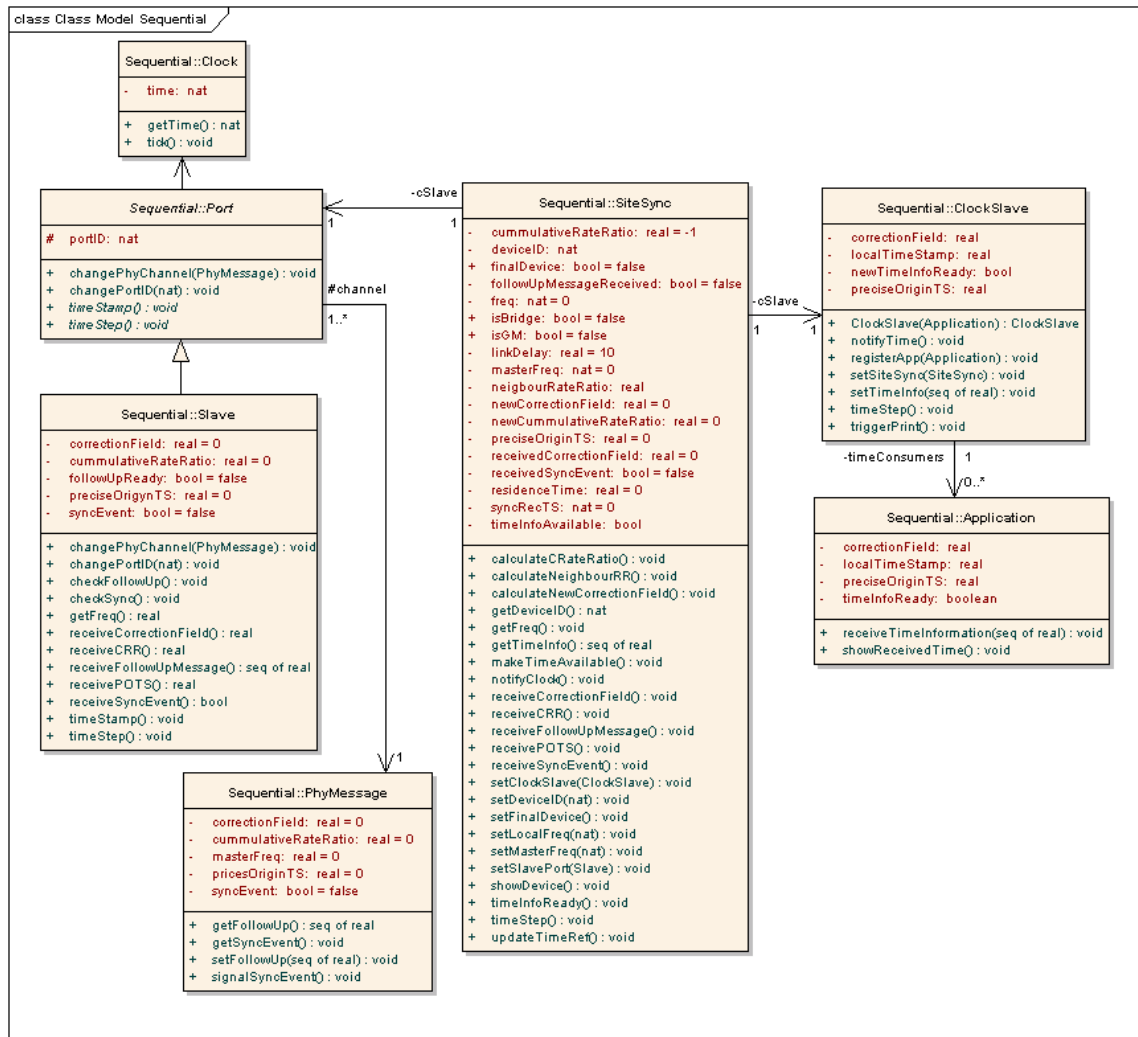


Figure 6.6: Sequential model UML class diagram.

timeStamp pattern has been applied. This pattern is implemented in the new class TimeStamp and used from the rest of the threads.

Figure 6.8 shows the class diagram of the concurrent version of the model. Attributes and methods have been omitted in order to increase readability. For additional details refer to figure 6.6. The stereotypes «thread» and «synchronized» are used to show active and concurrently accessed classes respectively.

6.4.3 Distributed real-time models

The distributed real time model gives the possibility of: A) Exploring different architectures and B) Analysing the real time deadlines accomplishment. Some changes are required before being able to carry out both studies. The first change has been the usage of the notion of time provided by Overture. The second change has been the addition of the Deployment class:

Deployment: defines the system infrastructure by the use of processing blocks (CPU) and communication interfaces (Buses). It contains the static class declarations of the modelled enti-

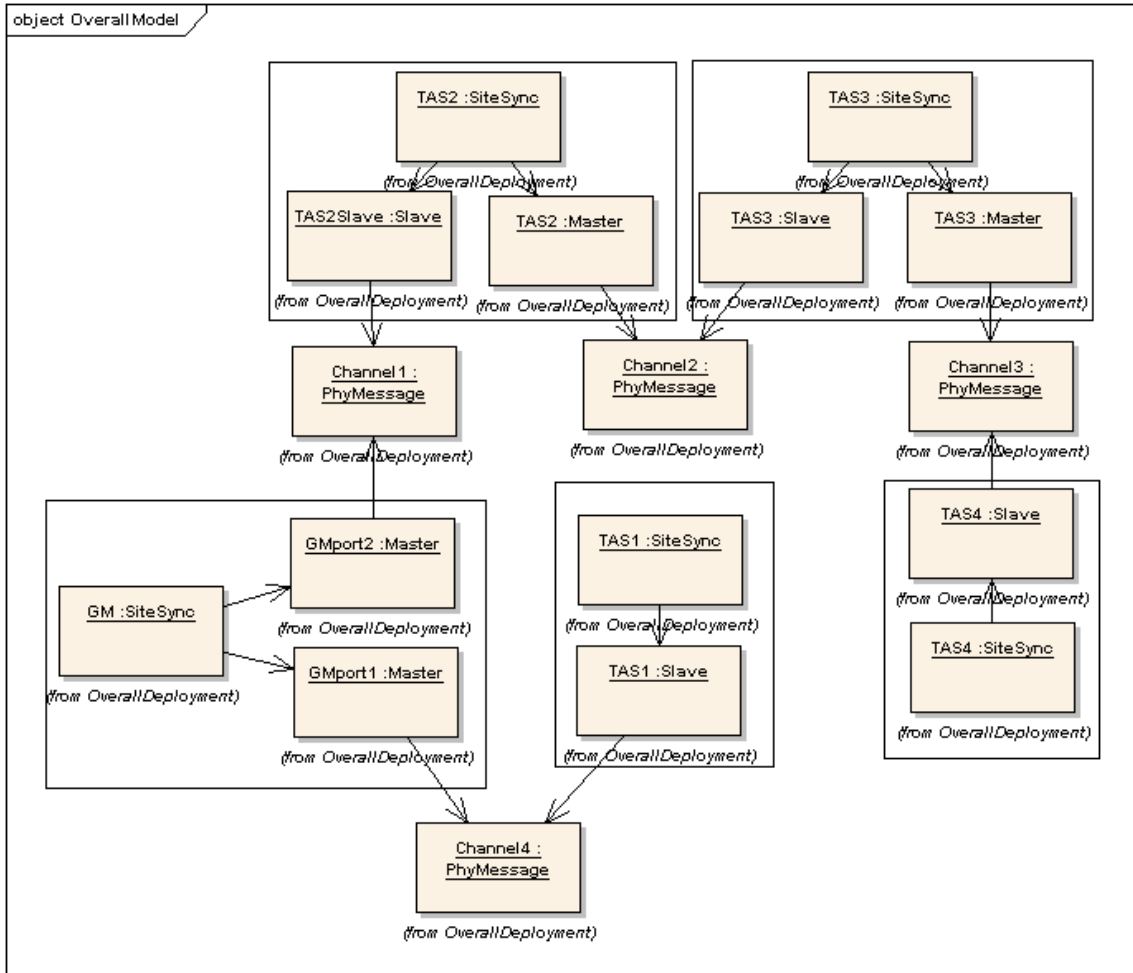


Figure 6.7: Context model UML object diagram.

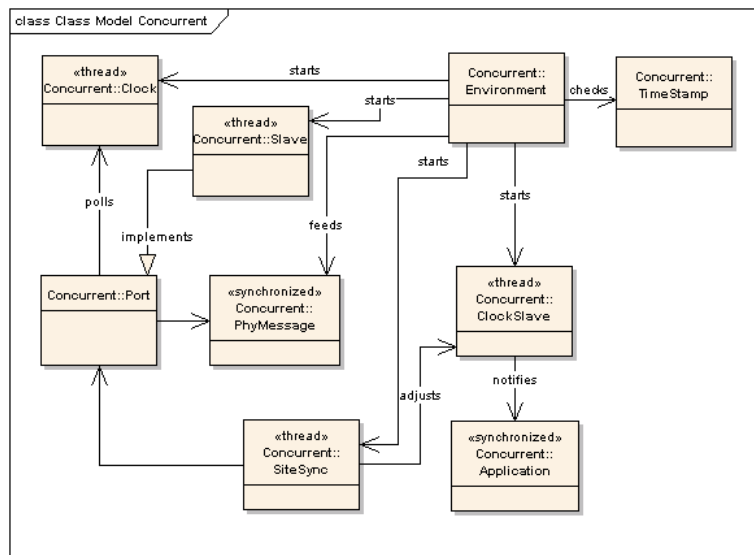


Figure 6.8: Concurrent model UML class diagram.

ties. This class is as well responsible for deploying the modelled entities in the processing units.

Figure 6.9 shows a class diagram modelling the main entities of the model and the existing relations among them.

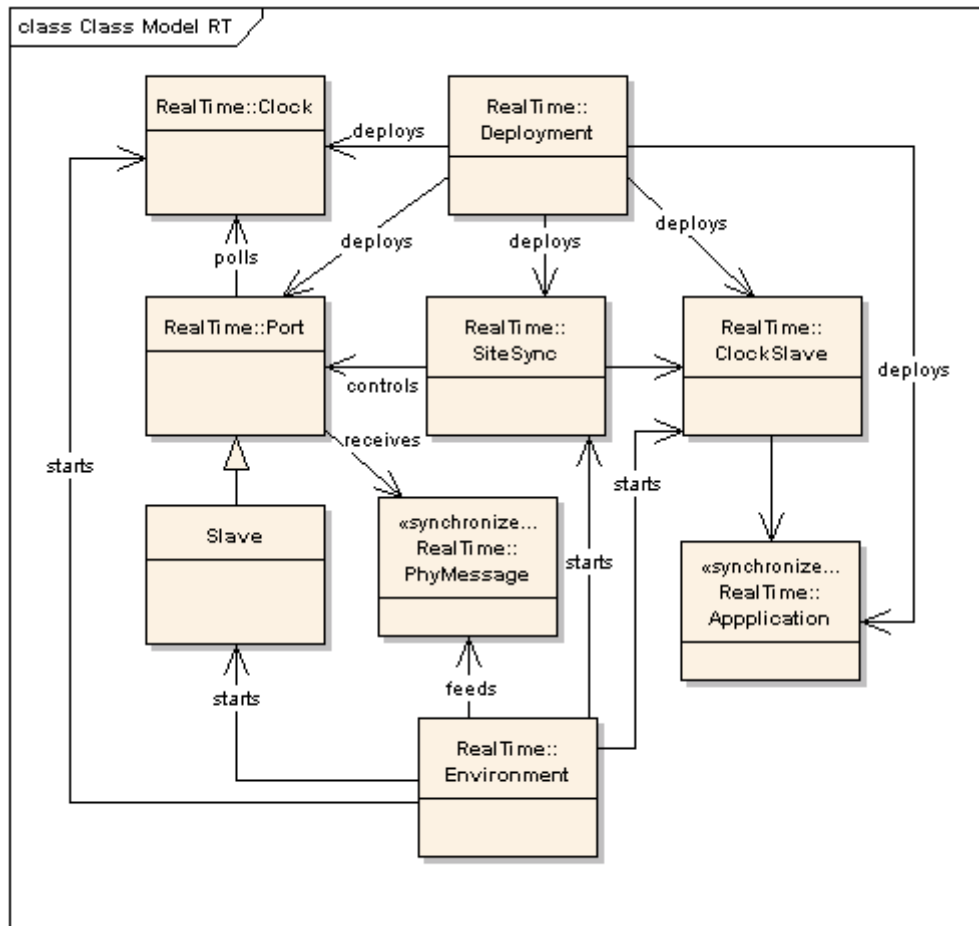


Figure 6.9: Real-Time model UML class diagram.

6.5. Evaluating different hardware/software architectures in a simulated deployment

Four different architectures have been evaluated during the exploration phase. In order to study the relevant timing aspects, log sequences have been created (application of guideline 12). In order to get the relevant values of time during model execution, the **time** instruction has been used (guideline 7).

6.5.1 Architecture 1: highly software based architecture

The first considered architecture is highly software based. Two CPUs have been used communicated through a hi-speed bus, modelling the available hi speed for hardware access. The ac-

cess to the channel is running on the CPU representing the hardware partition. This implies that the PhyMessage channel representing the hardware buffer for incoming time corrections will be deployed in hardware. A UML deployment diagram representing this architecture is shown in figure 6.10.

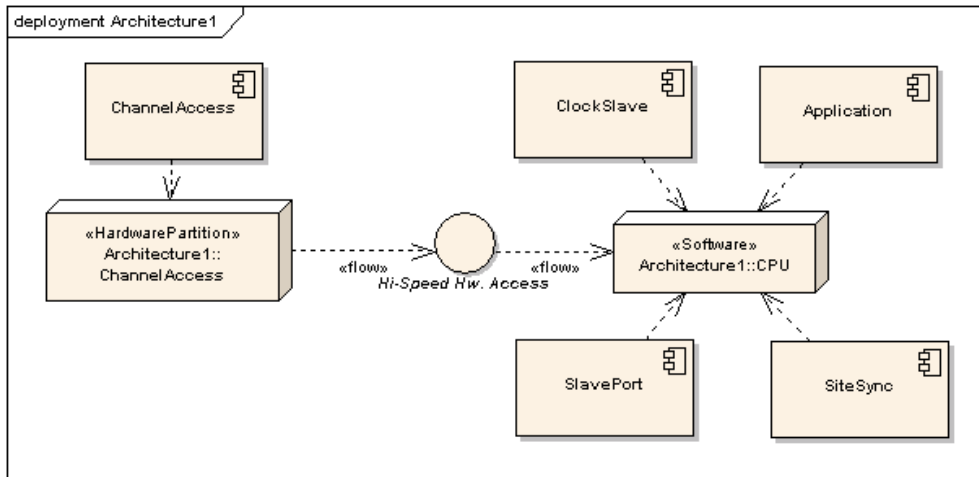


Figure 6.10: Deployment diagram for the highly software based architecture.

Model execution reveals that there is a considerable delay on the time stamping process under this architecture. This model has been analysed with CPUs of different frequencies. In the first case, a CPU running at a frequency of 1E7 Hz has been used. In this execution, the sync event was received at the hardware buffers 14ns after the simulation started. The same Sync event was timestamped after 4604 nanoseconds had elapsed. This constitutes a total delay of 4590 ns. In the second scenario a CPU running at a frequency of 10E6 Hz has been selected. The Sync event was still received at 14 ns, since the conditions of the environment did not change. On the other hand, the Sync event was time stamped after 46004 ns were elapsed, producing an error of 45990 ns. It can be conclude that such an implementation would be introducing an considerable delay of orders of magnitude in the synchronization process. Furthermore, this error would be depending on the speed of the used CPU and/or on its load at the moment of time in which the sync event was received.

Additionally, the maximum rate at which the endpoint device can receive time corrections is considerable reduced by this delay. If two sync events are received with a separation of 45 microseconds considering a cpu of 1Mhz, they will be timestamped at the same time therefore, the time information will not be valid².

6.5.2 Architecture 2: mostly software and offloaded SlavePort

The second architecture is making use of three different CPUs. The first CPU is modelling a hi-performance hardware partition. This hardware partition is communicated through a hi-speed bus to a second CPU. This intermediate is communicated to the final one through a low-speed bus. CPUs one and two are modelling regular general purpose processors. The channel access and the slave port have been deployed in the hardware partition. The synchronization logic SiteSync and ClockSlave is running in he intermediate CPU. Final time correction are made available to the third CPU through the Low-speed software interface. A UML deployment diagram representing this architecture is shown in figure 6.11.

²Unless this situation is properly identified and the second packet filtered out.

Architecture 3: offloaded SiteSync and SlavePort

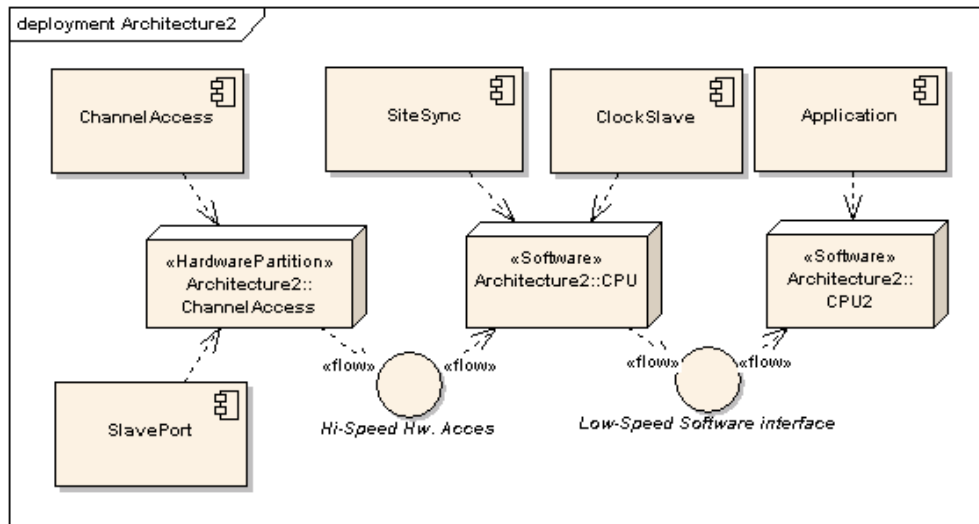


Figure 6.11: Deployment diagram for the mostly software with offloaded SlavePort architecture.

This architecture turns out to be much more accurate if compared with the first one. Model execution reveals that the introduced error in the time stamping is minimum. The model has been executed under two different configurations. The first configuration is the one shown in figure 6.11. In this case the Sync event is received after 28ns. The hardware is timestamping the sync event at 34ns. The total introduced error is 6ns. This result is very positive if compared with the one obtained by architecture 1. In order to explore a more cost-effective option, a second variant of architecture 2 is proposed. In this variant, the second CPU is not going to be used; the time consuming application is going to be deployed together with SiteSync and ClockSlave in the same CPU, running all of them as software components. The results are exactly the same as the ones obtained in architecture 2 configuration 1. This definitely supports the following thesis: A) The most critical and time sensitive process is the Sync event time-stamping, and B) The delay is considerably mitigated if the time-stamping process is carried out in the processing unit that is buffering the network traffic. The sync interval supported by architecture 2 is orders of magnitude above the one supported by architecture 1, and could be applied in very hi-speed transmission mediums such as optical fibres.

6.5.3 Architecture 3: offloaded SiteSync and SlavePort

The third architecture is keeping the access to the channel together with the slave port access in the hardware partition. The SiteSync logic has been allocated as well in the same hardware partition. A hi-speed bus is used to communicate the hardware partition with the second CPU. The second CPU is a general purpose processor running ClockSlave and Application as software components. A UML deployment diagram representing this architecture is shown in figure 6.12.

The purpose of this model is to explore how the endpoint device structure could be expanded to give support to more complex functionality, like the one that could be offered by a network bridge. In this case it has been decided to deploy the site sync entity in hardware. The prime focus of the model is to evaluate if there is a penalty in the time corrections reception while keeping the possibility of a more scalable design. The mode has shown that, the sync event is received 48 ns after the simulation has started, and this sync event is not timestamped until 74 ns have passed. This leads to a delay of 26 ns. The conclusions that can be drawn in this case are a little bit fuzzy. While it makes sense to consider hardware partitions as very high speed VDM-RT CPUs, it does

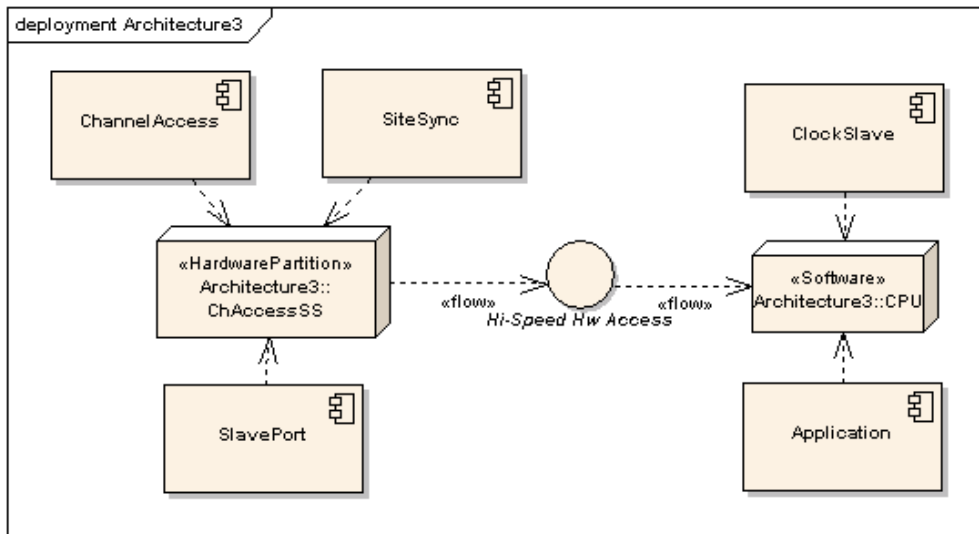


Figure 6.12: Deployment diagram for offloaded SiteSync and SlavePort architecture.

not make sense to use a single high speed CPU to deploy two potential hardware components on it. The reason behind is that active components are modelled as active classes, running as threads. These threads are going to be sharing computation time of the VDM-RT CPU. In the actual implementation, separate hardware partitions that can run in parallel will be implemented as purely separated hardware units, that will not be sharing "computation time".

6.5.4 Architecture 4: multiple hardware partitions

The fourth architecture is considering an improved hardware architecture. The goal is to improve the results obtained from architecture 3. In this case the first hardware partition is running the SlavePort and the ChannelAccess. This partition is communicated with a hi-speed bus with the second hardware partition, where SiteSync has been deployed. A second hi-speed bus has been used to communicate with the ClockSlave and the time consuming application. A UML deployment diagram representing this architecture is shown in figure 6.13.

This fourth architecture aims to model correctly the scenario proposed in architecture 3, a more scalable implementation with good time performance. In order to illustrate separate hardware units, two different hardware partitions have been used (application of guideline 9). Architecture 2 proved that the time stamping should be done at the hardware level, in the same block in which the buffers are located, so this design decision is going to be respected. Following the intent of architecture 3, site sync is going to be moved to hardware. Site sync is going to be able to access the hardware ports through the hi-speed hardware interface explained above, therefore the communication latency is cut down below 3 ns. Model execution showed that, the sync event is received at 28 ns and timestamped at 34 ns. This is keeping the 6 ns delay presented by the architecture 2. Site sync is able to retrieve the time corrections and execute the communication logic. In the case this architecture for the endpoint device would have to be used as a base to implement a compliant bridge, able to forward the time correction making use of master ports, this architecture would be correct. In the case the focus is to create a compliant endpoint device, this architecture would require additional development effort and very little positive impact (if any).

Model outcome assessment

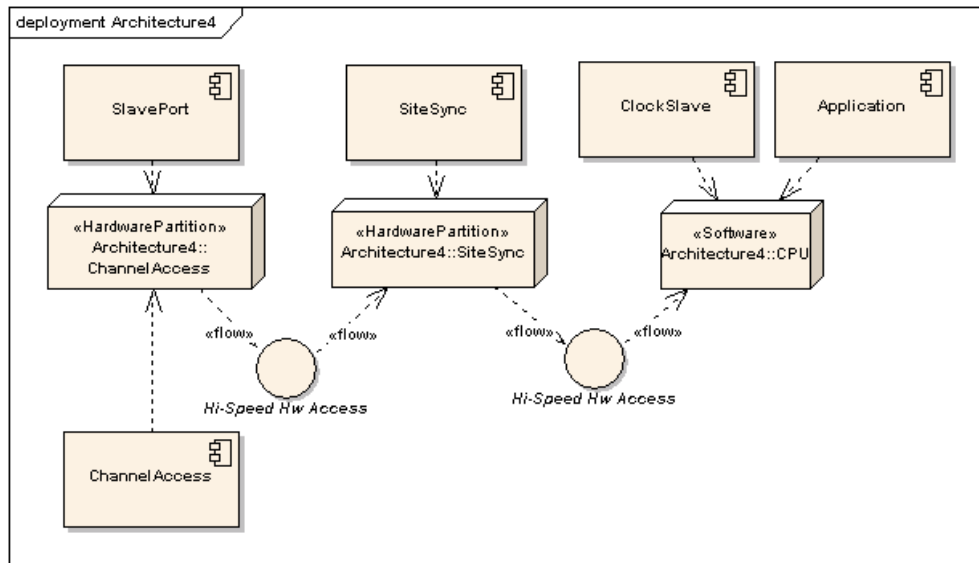


Figure 6.13: Deployment diagram for the multiple hardware partition architecture.

6.6. Model outcome assessment

The created models have been created much faster than the time it would have taking to produce partial prototypes of all four architectures before selecting the most appropriate for final implementation. The models have helped the system engineer to evaluate rapidly different architectures and, furthermore, to conduct analysis that will help him on deciding which factors will be improved among different architectures (such as scalability or possibility of further improvements). Besides the help on taking technical decisions, the application of modelling has bridged the mathematical description of the time synchronization protocol with a model that can be relatively easily understood by other professionals involved in system design.

Considering the time it takes to create a model like the one presented in this case study with regard the advantages it brings to the development process, the relation between invested resources/model outcome can be qualified, objectively, as very high.

6.7. Summary

This chapter has presented the application of the VDM-RT methodology for Hardware/Software co-design of embedded systems in a second case study, the development of an AVB endpoint device. This second case study is more complex than the introductory case study about servo control. The problem under study is more abstract as well, since the AVB case is at a system level, where the critical partitioning decisions have to be determined after initial study of the system. The proposed methodology have had a very positive impact on the problem understanding, helping the modeller on detecting the potential bottlenecks and critical aspects related to a problem unknown in advance. This has proved that the proposed approach not only performs well under "toy examples", but in real industrial cases as well. In the case this information would be presented to a systems engineer responsible for the AVB development project, he would be able to select which architectural candidate is the best one for the project needs. Furthermore, this selection would be based on solid information obtained by the application of a sound engineering approach.

Concluding remarks and future work

This chapter is composed of five sections. A brief review of the achieved results will be given in section 7.1. A more detailed assessment of the proposed methodology is presented in section 7.2. Future work and a review of further research lines in the application of the VDM-RT technology in Hardware/Software Co-design of embedded systems is presented in section 7.3. Reflections on the personal learning outcomes are presented in section 7.4 Final remarks are given in section 7.5.

7.1. Achieved results

The initial goals of this thesis, as presented in chapter 1 were:

- 1. To propose a general-purpose methodology based in the VDM-RT modelling language in order to support the design decisions in Hardware/Software co-design field.**
- 2. To propose improvements and extensions to the VDM-RT language and the modelling tool Overture, seeking to contribute on the creation of a modelling tool that can be applied in a more specialized field.**

Before being able to propose a new methodology for Hardware/Software co-design, it was important to study and review the state of the art in this field. Chapters 2 and 3 present a review of modelling languages and current approaches relevant for the development of Hardware/Software systems. Special emphasis was made on the System Level Design approach to Hardware/Software co-design. The modelling languages were evaluated with regard to the criterion presented by [Shaout&09], [Gajski&09], [Niemann98] and [Huang&04]. This technology and methodology review demonstrated that there is a need for a language and a methodology that enables the system engineer to perform: precise descriptions at a high abstraction level and allow rapid evaluation of different alternative architectural candidates in order to find the most appropriate hardware/software architecture.

The approach followed to accomplish the first goal is composed by two subgoals: the proposal of the methodology itself and the validation of the proposed methodology through its application in two case studies. In the three subsections below we assess our work in those two subgoals and the second goal of the thesis.

7.1.1 Development of a VDM-RT based methodology for Hardware/Software partitioning

Chapter 4 has proposed a new methodology that can support the hardware/software co-design process. The proposed methodology provides a reference for making partitioning decisions. These decisions are based on the accomplishment of real-time deadlines. Even though making partitioning decisions based on time constraints might be enough in some cases, there are other situations in which additional criterion shall be consider. As described in chapter 2, energy efficiency, delay at the hardware level or material costs are factors that should be taken into account in hardware/software co-design. The proposed methodology fails on considering these factors directly. This can be compensated by using other modelling technologies to extract information in those areas and consider it together with the architectures selected from the VDM-RT models. One of the strong points of the proposed methodology is the possibility of evaluating different hardware/software partitions with minimum changes on the models. On the other hand, the results provided by this exploration phase have to be used carefully and a certain tolerance factor should be used (additional details will be provided in section 7.2). Finally, the proposed methodology fails when approaching some localized partitioning problems. An example of a problem hard to tackle by this methodology is the signal processing field. This limitation is introduced by the language VDM-RT since it lacks of the necessary mathematical support in order to model signal processing algorithms. As it has been described above, the proposed approach in this case is to apply an heterogeneous modelling approach.

Additional assessment of this subgoal is presented in section 7.2.

7.1.2 Application in two case studies

The proposed methodology has been applied in two case studies with the purpose of validating it (presented in chapters 5 and 6). The case studies have shown that the application of the VDM-RT methodology for Hardware/Software co-design is providing relevant and sound information that can be used in the partitioning process. Case one, the development of a servo controller showed that this methodology can work on small cases. Predictions and results matched and it was shown that the methodology performed well under this scenario. Case two, the development of an AVB endpoint device shown that partitioning information can be extracted in complex problems as well. However, we cannot conclude that the partitioning decisions taken in that case were correct since we have not been able to produce the implementation for the four architectural candidates. Note that this would have implied producing and testing hundreds of thousands of VHDL code. The practical application of the methodology has shown that the usage of modelling in complex projects presents a high Return On Investment (ROI). The creation of models takes a considerable amount of time, but this initial investment can potentially be recovered later on in the development process. Additional assessment of this subgoal is presented in section 7.2

It can be concluded that we have accomplished goal one and reached the initial objectives set by its two subgoals. Once we proposed the methodology and after we applied it in the two case studies, we detected the need for concrete improvements for the VDM technology.

7.1.3 Proposal of improvements to the VDM technology

Chapter 4 has proposed a number of improvements to the VDM-RT language and the modelling tool Overture. Some of the proposed changes would be beneficial for making partitioning decisions with the proposed VDM-RT methodology.

The addition to *models of new platforms* would allow to conduct a more precise functionality performance evaluation. The concrete precision improvements would depend on the platform under study and they are complex to predict without studying particular cases. The *profiling of a model* can be already done by manually extracting the relevant information from the Overture log files. In some situations it is complicated to identify and to locate this information. The analysis that can be conducted with model profile information can lead to make partitioning decisions, hence the relevance of it. Modellers would have easier and clearer access to the information they need to make partitioning decisions if an structured and formatted execution report would be generated. Incorporating *annotations for real-time deadline evaluation* can have a very positive impact in the partitioning process. Some systems present complex time requirements that might depend on multiple factors. Trying to analyse all these factors by using log files every time the model is executed can be time consuming. Using annotations in order to describe time-critical aspects helps on formalizing the required time behaviour. Thanks to this formalization it is possible to automatically evaluate the real-time performance of the considered architectural candidates. An additional advantage is that this an step towards automatic design space exploration. An *improved bus support* would help on incorporating some relevant communication details. The improved bus support would rise the expressiveness of the models and it would simplify the communications modelling by eliminating the proxy pattern. Additional assessment on communication modelling is provided in section 7.2.

Automatic variable monitoring would make it easier to analyse the evolution of a certain variable over time. Like annotations for real-time deadline evaluation, this feature would be a step forward towards automatic design space exploration. Automatic variable monitoring would make the modelling and analysis process less complex but it would not improve the quality of the partitioning decisions.

The addition of *code generation* is complex to implement in an efficient manner. Considering that this methodology is targeting the development of embedded systems optimization with regard to speed and code size is a requirement. While mapping VDM++ models to C++ code (or another Object Oriented language) is something that could be explored and partially implemented in a Master's thesis, doing the same with VHDL would be more difficult. The reason resides in the existing abstraction gap between VDM and VHDL. Even though code generation would allow a smooth progression from the modelling to the implementation stage, it would have very little positive impact (if any) on making partitioning decisions from VDM-RT models. This makes code generation optional in this context.

The proposed improvements are worthwhile considering from the user perspective. Incorporating models of new platforms and improving the bus support would allow the user to extract more precise and reliable information, gaining confidence in the methodology proposed. The incorporation of annotations for real-time deadline evaluation would give the user the possibility of carrying out more complex analysis in a faster way, increasing his performance considerably. Finally, model profiling and automatic variable monitoring would make the methodology more usable and it would eliminate the need for external scripts and tools.

It can be concluded that we have accomplished goal two, not only by just presenting a features wishlist but by describing a set of improvements, the rationale for its consideration and how this methodology can benefit from them. Additional assessment of this subgoal is presented in section 7.2

7.2. Assessment of the VDM-RT methodology for Hardware/Software co-design

This thesis has made extensive use of abstract models in order to represent the combination of hardware and software in the development of systems. More precisely, the models have been used to extract information to support the hardware/software partitioning process. During the process of creating these abstract models, some details have not been considered, since they have been left out intentionally by the modeller. Other details have been left out by the modelling technology itself. Being aware of what has been lost during this formalization is required in order to move forward in the development process, bridging the gap between modelling and implementation. A weakness of this methodology is that how to bridge this gap might not be obvious in all cases. Some of these problematic situations are introduced by the notion of time used in VDM-RT, the VDM-RT kernel and the modelling of communication.

7.2.1 The notion of time

The VDM-RT language assumes that every CPU deployed in a model is perfectly synchronized with regard to the rest. Considering a single FPGA system, such an assumption can be accepted. In the case the problem under study is a distributed real-time system, it might be necessary to consider additional synchronization mechanisms in the implementation stage. Some implementation technologies provide abstraction layers taking care of the time synchronization so, depending on the used technology, engineers might not need to produce the time synchronization implementation. Note that these synchronization problems cannot be captured by the VDM-RT models and therefore this is a case in which the new methodology does not support the systems engineer. This is a limitation introduced by the modelling technology since VDM-RT is targeting the "big picture" of the system rather than time synchronization in a distributed environment. Further details and mechanisms to compensate this problem are described in [Babaoglu&93].

7.2.2 Time precision

A word of caution has to be given about the precision of the time predictions offered by the VDM-RT models. As it has been described above, the VDM-RT system models are coarse-grained. The data these models are providing for analysis is good enough to take co-design decisions by comparison. If this comparison is performed between figures provided by the model, no additional precautions are needed. In the case this comparison is performed between figures provided by the abstract model and figures provided from the real world, a tolerance of at least one order of magnitude should be introduced. This might be imprecise in some situations, but these models are created at a high level and are not targeting very fine-tuned implementation details. As described in section 3.3, the System C modelling library incorporates a notion of time that can reach the femtosecond level, orders of magnitude below the nanosecond level reached in VDM-RT. Furthermore, this notion of time is directly connected with the notion of time used in VHDL. Making use of SystemC can be a good alternative to reach a higher degree of precision in the models. The price to pay for this precision: less abstract and complex representations if compared to VDM-RT models.

7.2.3 The RTOS abstraction

The VDM-RT kernel is modelling a Real-Time Operating System layer. This layer provides the mechanisms to manage different types of threads and synchronization predicates like mutexes and

history counters. These mechanisms are implemented in many different ways by RTOS in the market. VDM-RT uses a different paradigm for synchronization and protection of critical sections. For example, the functionality offered by history counters in VDM-RT might not be explicitly available in a certain RTOS, but surely a different mechanism will be provided in order to accomplish the same task. This relation holds in the other way. Some of the features provided by RTOS implementations are not present in VDM-RT by default. For instance, critical sections are present in many RTOS implementations but not in VDM-RT. Mechanisms like counting semaphores have to be explicitly modelled in VDM-RT if they are going to be used in the models, while they are available by default in RTOS implementations. The system engineer has to be aware about the possible implementation options offered by different technologies while mapping a model to a real implementation. Besides these implementation details, time performance of different RTOS implementations needs to be taken into account. Under different RTOS it might take a different amount of time to perform a context switch, to spawn a thread or to acquire a lock. VDM-RT is introducing its own timing characteristics in those aspects. Again, these differences have to be carefully evaluated when moving to the final implementation.

7.2.4 Modelling communication

Communication between systems and subsystems is one of the problematic issues in system design. The suggested methodology is focused on the functionalities, more precisely on the real-time performance of the system functionalities. In order to evaluate the functionality the communication can be partially modelled. The proposed approach could be used to determine the time-frame and approximate boundaries for the communication facilities required by the modelled functionality. If the problem under study is focused on the communication performance, this methodology will most likely be not appropriate to address it. This shortcoming is introduced by the representation of the communication link by BUSES, which might not be precise enough. Some high-level aspects need to be incorporated in the BUS representation, like the bus load or possible traffic losses. Other aspects are very difficult to capture in VDM and therefore it does not make sense to consider them. Communication depends heavily on physical phenomena, like noise, temperature, distance or wireless propagation. These factors are entirely abstracted away in VDM-RT, even though they are relevant in communication modelling, it is difficult to analyse them from the proposed methodology. As seen in [Nielsen10b] VDM-RT as it is at the moment is not ideal to model dynamic architectures with wireless devices. As a conclusion regarding the proposed methodology and the communication problem, the modeller must be aware of the root purpose of the analysis conducted on the model, the consequences and the final impact in the system functionality. By doing so, many communication aspects can be considered with regard to the created models even though they have not been fully detailed.

7.2.5 Code generation

Code generation has been described as one of the desirable tool features from which this methodology could benefit. Considering that the code is generated from models that are not depicting the whole system in detail, it is not possible to obtain a final implementation automatically. The problematic details discussed above regarding time, concrete RTOS behaviour and communication would have to be very carefully studied, making sure that the assumptions taken are respected at the implementation level. It is important to keep in mind that the final product is going to be implemented by using technologies different from the paradigms used in the models. The differences between the implementation and the modelling technologies, have to be considered in order to overcome the gap between the abstract model and the real world.

7.2.6 Application in two case studies

The new suggested methodology has only been applied in two case studies, and compared to final implementations in only one. The results look promising but, to determine if such a methodology would work in general, it should be applied in more cases (preferably of different nature) and have other people to test it. This assessment cannot be carried out in one MSc. thesis, so therefore we can conclude that the results of the suggested methodology look promising but a wider analysis in terms of different projects, professionals and longer time is needed to determine the limits of its capabilities.

The servo controller case was conceptually and technically simple, and it made it possible to show in practice how the methodology could be applied. Thanks to this simplicity it was possible to go down to the implementation level and compare the predictions with actual implementation results. The AVB case was much more challenging than the servo case from a technical point of view. The protocols used in AVB were complex and reaching a good understanding was not an easy task. The application of models helped us on going through these problems, since they forced us to be systematic and to evaluate which details were worth considering in the partitioning problem. The models created made it possible to select among several different architectural candidates without having to perform an actual implementation. Due to time limitations an implementation of the AVB endpoint was not carried out. Like in the servo case, an implementation would have made it possible to relate predictions with actual results. However, an AVB implementation can take up to several VHDL thousands of LOC and is not feasible for a Master's thesis of 6 man-month. It is difficult to provide an accurate estimate of the time it would take to produce functional implementations for the four architectural candidates we have considered in this work. The development time can be shortened if we implement a limited architecture aiming to validate the models. These limited implementations, even though not fully functional AVB devices, could provide a preliminary structure for further development. The real-time features studied in the models could be implemented and deployed in these preliminary hardware/software stubs. If we follow this strategy, we could create the four implementations and support the model conclusions. This kind of implementation considered alone is expected to have a workload comparable to this MSc. thesis project.

Considering the size of the models, the servo case has been modelled with only 500 LOC in VDM-RT. On the other hand, the final implementations have required 260 LOC in C for the software only solution and 1195 LOC in VHDL and 382 LOC in C for the hardware/software solution. This shows that only 500 lines of VDM have been enough to evaluate architectural implementations with a total size of 1738 LOC of varying complexity. In the AVB case study the size of the VDM-RT model is 1083 LOC. This model is presenting how the time synchronization protocol is handled by the device in four different architectures, some of them using different configurations. The AVB protocol implementation provided by Xilinx¹ is composed by 55000 VHDL LOC. Note that the architectural pillars of this solution have been tackled from a VDM-RT model of 1083 LOC. Considering the complexity of the protocol and the number of evaluated architectures it would be fair to conclude that it is an expressive model. It would be difficult to extract the same kind of information from a SystemC, SysML or Matlab/Simulink model of similar size.

Finally, it is worth to remark that the servo case models were created at a much lower level of abstraction than the AVB protocol models. The servo case was focused on a very concrete opera-

¹Xilinx produces an IP core implementing the AVB protocol. Additional information can be found in [XilinxAVB11]

tional detail. This explains why there is such a small difference in size between VDM-RT models and final implementation.

7.2.7 Comparison with other methodologies

Chapter 3 presented different methodologies for the development of Hardware/Software systems. It is worth to assess how the proposed methodology performs in comparison to them. The methodology proposed by Xilinx is at an operation level and can hardly cope with complex partitioning decisions. The methodology based on Matlab/Simulink is strong in modelling problems involving heavy mathematical computations, for example signal processing problems. However, it fails when considering problems that would benefit from being described at a higher level of abstraction. The Saturn methodology, a combined SysML-SystemC approach, exploits co-simulation possibilities bridging the gap between modelling and implementation. However, it might be too close to the implementation level. The SHE methodology is close to the methodology proposed in this thesis. The SHE methodology is using the modelling language POOSL instead of VDM-RT. The expressive power of POOSL is comparable to VDM-RT and it presents many other similarities in terms of constructs. Additional effort has been spent in the development of predictable control software able to cope with real-time deadlines, a feature that is missing in VDM-RT. There has been no new information in the POOSL/SHE area since year 2007 and the state of the project is unknown. Even though the SHE/POOSL methodology is interesting, taking into account its current development state, it can only be considered as a good source of inspiration for further developments in VDM-RT.

Compared to the Matlab/Simulink approach, the methodology proposed in this thesis is more abstract and far from concrete implementation technologies. However, it cannot offer the same performance on problems involving complex mathematics. The Saturn approach is better in virtual prototyping, but limited to concrete platforms. In comparison, the proposed methodology is platform independent. The VDM-RT based methodology is strong in high level design space exploration but further effort is required to relate the design alternatives considered at this level with the technologies uses in the implementation phase.

7.3. Future Work

This thesis work has given the possibility of detecting interesting areas and concrete action points that could be explored in further developments. Below suggestions for these are provided.

7.3.1 Implementation of the AVB endpoint device proposed architectures

In chapter 6 four different architectural candidates were modelled. These models provided information relevant for the partitioning process and helped on selecting the optimal candidate for implementation. Due to time constraints and complexity of the problem, no early prototype generation was performed. Possible further work could be the implementation of the four architectures discussed. Once the implementation would be complete, it would be interesting to validate the information extracted from the VDM-RT models against the prototypes created. This would reinforce the validity of the proposed methodology, by demonstrating its applicability in a complex case from the design phase down to the implementation level.

7.3.2 Guidelines for the transformation of VDM-RT to SystemC

SystemC is one of the most used languages in the modelling and design of hardware/software systems. There are tools available in the market that allow the generation of C/C++ code for software components and VHDL code for hardware components. SystemC could act as a bridge between abstract VDM models and the final implementation.

Besides partially solving the of lack of VHDL code generation from the VDM modelling environments, SystemC will bring a stepwise evolution towards the final implementation. This would result in an intermediate modelling stage in SystemC, that will allow to perform modelling at a lower abstraction level. Such a low level model would possibly enable the incorporation of more low level details in the case they are needed for more concrete architectural evaluation. On the other hand this level would be less abstract and would require more modelling time. Aspects that were not detailed at the VDM level would have to be incorporated at the SystemC level. This would require additional modelling time and the complexity of the representation of the system will increase.

A systematic approach would be required in order to move the VDM models to SystemC models. The proposal of such an approach, and the study of the equivalences between the represented entities in VDM and the SystemC constructs should be studied in detail in further work.

7.3.3 Exploiting further the possibilities offered by Overture and VDM-Tools

Both Overture and VDM-Tools present features that make creating higher quality models easier. These possibilities were not explored during the development of the case studies. Some of these features are combinatorial testing, proof obligations and model coverage. The application of combinatorial testing allows to evaluate the created models against all possible combinations of input values, performing a thorough testing of the model. By ensuring that the model is performing properly under all kind of inputs from the environment we are reducing modelling errors. The application of proof obligations make sure that the error sources introduced by the way the model has been created are fixed. This improves the overall robustness of the model. Finally, model coverage indicates to which extent the model has been evaluated. A coverage of 100% is obtained in case all the parts of the model have been executed at least once. It is desirable to have the highest percentage of coverage possible, implying that most of the model has been covered under different runs. By verifying and testing the model by the means exposed above, confidence in the used method for taking the partitioning decisions is increased.

7.3.4 Applying a heterogeneous modelling approach based on the VDM-RT language, 20-sim and SysML

As discussed in this thesis, only a combination of several modelling technologies can produce a complete description of the system. In the same way, a complete analysis of the system can only be achieved by using multiple modelling technologies. Following this idea, the proposed extension to this work is the incorporation of additional information in terms of interaction with external systems and physical world. New modelling languages should be used in order to described these specialized areas, the languages SysML and 20-sim. A link between these modelling languages and the VDM methodology would be required in order to fully exploit the possibilities of this heterogeneous approach.

The projects [COMPASS11] and [DESTTECS09] are providing tools linking SysML and 20-sim with the modelling language VDM. Additionally, these projects incorporate specific facilities to conduct analysis in the logical and physical world. Such an analysis would be very difficult

Personal Learning Outcomes

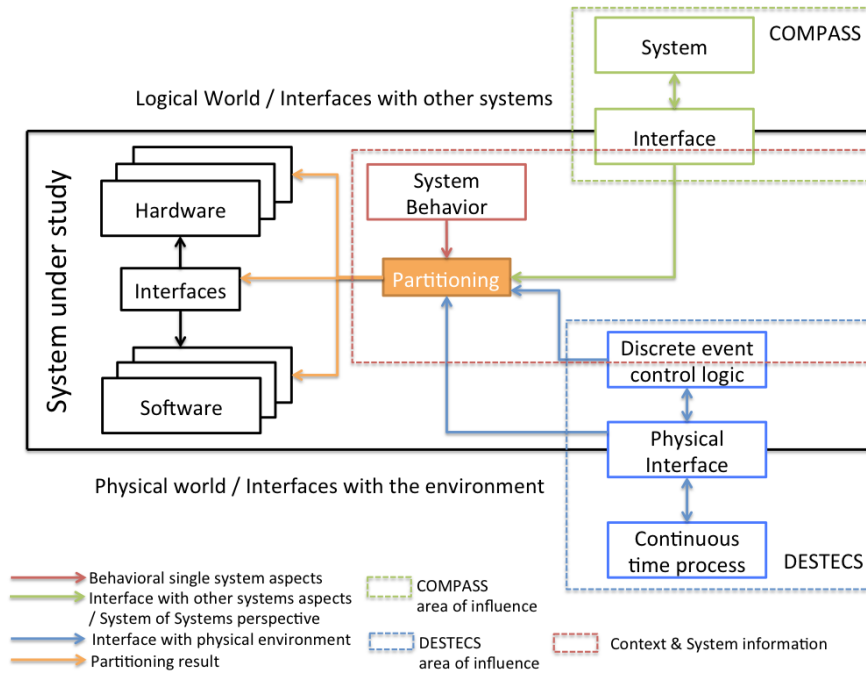


Figure 7.1: Heterogeneous modelling approach.

to be achieved by a VDM-only approach. Integrating the models created under COMPASS and DESTECS in the partitioning process would relate the implications of certain system architectures to the environment in which the system would be operating.

Figure 7.1 shows the proposed relation between the COMPASS and DESTECS projects and the partitioning process.

7.4. Personal Learning Outcomes

As described above, one of the case studies was based on the Audio Video Bridging protocol. Before modelling an AVB compliant device it was a necessary to understand the protocols behind it and the hardware and software internals of an AVB device. This was specially complicated and it took me weeks of reading, meetings, and several e-mails with hardware manufacturers. AVB has been specified in 2011 therefore the documentation available aside from the standards is reduced to a few academic papers. One of the mistakes I made while trying to understand the protocols behind AVB was to start with the 450 pages standard, trying to get every single detail on it. After two weeks following this track I decided to read the few available papers about AVB instead and revisit the standard right after that. This turned out to be a good approach. The papers gave me the overall picture I needed to relate the details presented in the standard. Surprisingly, this kind of "overall picture" was not part of the standard that was focused on operational details. The papers allowed me to start creating an overall VDM model of the system. I continued expanding this initial model with the relevant details extracted from the standard documentation. As a result of this modelling stage, I personally think that the executable VDM model I have created expresses better how AVB works than 400 pages of text. The conclusion that I can extract from this experience is that the next time I will face the problem of understanding a complex system or protocol I will read the scientific papers before delving into the standard documentation.

FPGAs are the most relevant platforms in the Hardware/Software Co-design field. My experience with this technology was very limited when I started this MSc. thesis. Conducting this work gave me the opportunity to learn more about them, mostly by reading papers, doing small experiments with FPGAs (not included in this thesis work since they are not related to it) and following a 4-day course organized by the Xilinx FPGA manufacturer. Even though I am not an expert on FPGAs I can say that I have a better understanding of its internals, possibilities and limitations. This gives me the basic knowledge to judge if they are a good technological alternative in the case I have to design an embedded system in the future.

The competences acquired in the courses "Model-Driven Development" and "Hardware/Software co-design" have been complemented with more in-depth new learning outcomes. In the model-driven development course we were introduced to modelling and the VDM language, but we did not analyze in detail the information obtained from models. One of the main pillars of this thesis has been to study and extract information from models, therefore, I have improved my skills and understanding of how modelling can be used in practical cases. In the Hardware/Software co-design course we were introduced to informal models, graphical notations and methodologies for Hardware/Software systems development. In this thesis I have reviewed these modelling technologies and methodologies, compared them and identified their strong and weak points. This has given me the opportunity to exercise critical thinking and analytical skills in a technical context.

Additionally, the process of creating this thesis has been an starting point in the process of becoming better in writing and communicating technical ideas. After finishing this thesis I see that one of the most complex parts of writing is being selective, analysing the contribution of part of the text to the global argument presented by the thesis. This implies to decide whether certain information has to be written as part of the main body, incorporated in appendixes or only cited.

It is the first time during my academic life in which I have had to create an argument of this size and support it properly. If I compare this MSc. thesis with my previous work I can clearly see a difference between them. My Bachelor's thesis was very product-oriented, therefore its report was a descriptive document focused on technical details. During my Master's degree some of the courses I have followed have required the creation of projects with an additional level of judgement, comparison and reflection. I now feel for the first time that I understand what doing research involves. Finally, this thesis has required a more mature approach, better planning and organization and to relate my work with current research work.

7.5. Final Remarks

This thesis has successfully approached the initial objective of developing a VDM-RT based methodology that can help on taking partitioning decisions in the Hardware/Software Co-design of embedded systems. This initial objective was a challenge, given the abstraction gap between a VDM representation and a final implementation involving the development of hardware. The development of the proposed methodology has been justified and compared against the existing approaches, showing that there is a place for a VDM-RT based approach in the development of Hardware/Software systems.

The application of the proposed methodology in two case studies has shown that the presented thesis can provide solid arguments for a particular candidate architecture. The proposed additions

Final Remarks

to the tools and the language can improve the performance of the proposed VDM-RT methodology considerably for hardware/software co-design.

Some of the additions could be developed under new student projects or MSc. thesis projects. The application of formal methods and modelling in the Hardware/Software Co-design field is a topic that still needs further research and development. This thesis has opened a new way in which this research could continue.

References

- [Amos&02] S. W. Amos and R. S. Amos. *Newnes Dictionary of Electronics*. Newnes, fourth edition, 2002. 389 pages. [cited at p. 113, 114, 115]
- [Babaoglu&93] Özalp Babaoglu and Keith Marzullo. *Consistent Global States of Distributed Systems: Fundamental Concepts and Mechanisms*. Technical Report UBLCS-93-1, University of Bologna, Piazza di Porta S. Donato, 5, 40127 Bologna (Italy), January 1993. [cited at p. 94]
- [Black&04] David C. Black and Jack Donovan. *SystemC: From The Gourd Up*. Springer, first edition, 2004. 244 pages. [cited at p. ix, 1, 6, 10, 19, 20, 26, 31]
- [Broenink&10] Broenink, J. F. and Larsen, P. G. and Verhoef, M. and Kleijn, C. and Jovanovic, D. and Pierce, K. and Wouters F. Design Support and Tooling for Dependable Embedded Control Software. In *Proceedings of Serene 2010 International Workshop on Software Engineering for Resilient Systems*, pages 77–82, ACM, April 2010. [cited at p. 70]
- [Brogioli&06] Michael Brogioli and Predrag Tadosavljevic and Joseph R. Carvallaro. Hardware/Software Co-design Methodology and DSP/FPGA Partitioning: A Case Study for Meeting Real-Time Processing Deadlines in 3.5G Mobile Receivers. 2006. 5 pages. [cited at p. 17]
- [Brooks86] Frederick P. Brooks Jr. No silver bullet - essence and accidents of software engineering. *Information Processing*, 7, 1986. 1069-1075 pages. [cited at p. 6]
- [Bruel&98] Jean-Michel Bruel. Integrating Formal and Informal Specification Techniques. Why? How? In *Proceedings of the Second IEEE Workshop on Industrial Strength Formal Specification Techniques*, pages 50–, IEEE Computer Society, Washington, DC, USA, 1998. [cited at p. 24, 28]

Chapter 7. Concluding remarks and future work

- [Burns&09] Alan Burns and Andy Wellings. *Real-Time Systems and Programming Languages*. Ada, Real-Time Java and C/Real-Time POSIX. Addison-Wesley, fourth edition, 2009. 602 pages. [cited at p. 31]
- [Chen&09] Sao-Jie Chen and Guang-Huei Lin and Pao-Ann Hsiung and Yu-Hen Hu. *Hardware Software Co-Design of a Multimedia SOC Platform*. Springer, 2009. 151 pages. [cited at p. 16, 19, 26]
- [Cofer&06] R.C. Cofer and Benjamin F. Harding. *Rapid System Prototyping with FPGAs: Accelerating the design process*. Embedded Technology, Elsevier Inc., Linacre House, Jordan Hill, Oxford OX2 8DP, UK, 2006. 301 pages. [cited at p. 114, 115, 116]
- [COMPASS11] Comprehensive Modelling for Advanced Systems of Systems. 2011. <http://www.compass-research.eu/>. [cited at p. 5, 98]
- [DESTECS09] DESTECS (Design Support and Tooling for Embedded Control Software). European Research Project, June 2009. <http://destecs.org>. [cited at p. 5, 98]
- [Edwards&97] Stephen Edwards and Luciano Lavagno and Edward A. Lee and Alverto Sangiovanni-Vicentelli. Design of Embedded Systems: Formal Models, Validation and Synthesis. *Proceedings of the IEEE*, 85(3):366–390, March 1997. 22 pages. . [cited at p. 255]
- [Fitzgerald&05] John Fitzgerald and Peter Gorm Larsen and Paul Mukherjee and Nico Plat and Marcel Verhoef. *Validated Designs for Object-oriented Systems*. Springer, New York, 2005. [cited at p. 30]
- [Gajski&00] D. D. Gajski and J. Zhu and A. Gerstlauer and S. Zhao. *SpecC, Specification Language and Design Methodology*. Springer, 2000. 336 pages. [cited at p. 13]
- [Gajski&09] Daniel D. Gajski and Samar Abdi and Andreas Gerstlauer and Gunar Schirner. *Embedded System Design, Modelling Synthesis and Verification*. Springer, 2009. 347 pages. [cited at p. ix, 1, 10, 11, 12, 19, 20, 23, 24, 26, 91]
- [Garner&07] Garner, G.M. and Feifei Feng and den Hollander, K. and Hongkyu Jeong and Byungsuk Kim and Byoung-Joon Lee and Tae-Chul Jung and Jinoou Joung. IEEE 802.1 AVB and Its Application in Carrier-Grade Ethernet [Standards Topics]. *Communications Magazine*, IEEE, 45(12):126–134, December 2007. 8 pages. [cited at p. x,

Final Remarks

74, 75]

- [Garner&09] Garner, G.M. and Gelter, A. and Teener, M.J. New simulation and test results for IEEE 802.1AS timing performance. In , editor, *Precision Clock Synchronization for Measurement, Control and Communication*, 2009. ISPCS 2009. International Symposium on, pages 1–7, IEEE, IEEE, Los Alamitos, CA, USA, October 2009. 7 pages. [cited at p. 80]
- [Garner&11] Garner, G.M. and Hyunsurk Ryu. Synchronization of audio/video bridging networks using IEEE 802.1AS. *Communications Magazine*, IEEE, 49(2):140–147, February 2011. 7 pages. . [cited at p. 80]
- [Huang&04] Jinfeng Huang and Jeroen Voeten and Andre Ventevogel and Leo van Bokhoven. *Platform-independent Design for Embedded Real-Time Systems*, chapter 1, pages 35–50. Kluwer Academic Publishers, Norwell, MA, USA, 2004. [cited at p. 4, 14, 23, 91]
- [IEEE-std1722-2011] The Institute of Electrical and Electronics Engineers, Inc. *IEEE Standard for Layer 2 Transport Protocol for Time-Sensitive Applications in Bridged Local Area Networks*. Technical Report, IEEE, NY, USA, May 2011. 45 pages. [cited at p. x, 73, 74]
- [IEEE-Std802.1As-2011] The Institute of Electrical and Electronics Engineers, Inc. *IEEE Standards for Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks – Local and Metropolitan Area Networks*. Technical Report, IEEE, NY, USA, 2011. 274 pages. [cited at p. x, 75, 76, 77, 80]
- [IEEE-Std802.1Qat-2010] The Institute of Electrical and Electronics Engineers, Inc. *IEEE Standards for Virtual Bridged Area Networks. Amendment 14: Stream Reservation Protocol (SRP) – Local and Metropolitan Area Networks*. Technical Report, IEEE, NY, USA, 2010. 101 pages. [cited at p. 75]
- [IEEE-Std802.1Qav-2009] The Institute of Electrical and Electronics Engineers, Inc. *IEEE Standards for Virtual Bridged Area Networks. Amendment 12: Forwarding and Queuing Enhancements for Time-Sensitive Streams – Local and Metropolitan Area Networks*. Technical Report, IEEE, NY, USA, 2009. 71 pages. [cited at p. 74]
- [INCOSEseh10] International Council on Systems Engineering. *Systems Engineering Handbook. A guide for system life cycle processes and activities*. Technical Report, INCOSE, 7670 Opportunity Rf, Suite 220 San Diego, CA, January 2010. [cited at p. 9]

Chapter 7. Concluding remarks and future work

- [ISOVDM96short] P. G. Larsen and B. S. Hansen and others. Information technology – Programming languages, their environments and system software interfaces – Vienna Development Method – Specification Language – Part 1: Base language. December 1996. International Standard ISO/IEC 13817-1. [cited at p. 29]
- [Kramer07] Jeff Kramer. Is Abstraction the Key to Computing? *Communications of the ACM*, 50(4):37–42, 2007. [cited at p. 10]
- [Larsen&09] Peter Gorm Larsen and John Fitzgerald and Sune Wolff. Methods for the Development of Distributed Real-Time Embedded Systems using VDM. *Intl. Journal of Software and Informatics*, 3(2-3), October 2009. [cited at p. 4, 5, 41]
- [Larsen&10a] Peter Gorm Larsen, Sune Wolff, Nick Battle, John Fitzgerald and Ken Pierce. Development Process of Distributed Embedded Systems using VDM. Technical Report TR-2010-02, The Overture Open Source Initiative, April 2010. [cited at p. ix, 37, 39]
- [Larsen&10b] Peter Gorm Larsen and Nick Battle and Miguel Ferreira and John Fitzgerald and Kenneth Lausdahl and Marcel Verhoef. The Overture Initiative – Integrating Tools for VDM. *ACM Software Engineering Notes*, 35(1):, January 2010. 6 pages. [cited at p. 5, 30]
- [Larsen&10c] Peter Gorm Larsen and Kenneth Lausdahl and Nick Battle. *The VDM-10 Language Manual*. Technical Report TR-2010-06, The Overture Open Source Initiative, April 2010. [cited at p. 30]
- [Lausdahl&10b] Kenneth Lausdahl and Peter Gorm Larsen. VDM-RT Scheduling in Overture. April 2009. 16 pages. Model driven development using VDM++ and UML II. [cited at p. ix, 4, 5, 6, 32, 33, 43, 55]
- [Matlab] MathWorks. <http://www.mathworks.com>. October 2011. Matlab official website. [cited at p. 27]
- [Mischkalla&10] Mischkalla, Fabian and He, Da and Mueller, Wolfgang. Closing the gap between UML-based modeling, simulation and synthesis of combined HW/SW systems. In *Proceedings of the Conference on Design, Automation and Test in Europe*, pages 1201–1206, European Design and Automation Association, 3001 Leuven, Belgium, Belgium, 2010. 6 pages. [cited at p. 5, 35]
- [Mueller&10] Mueller, W. and Da He and Mischkalla, F. and Wegele, A. and Whiston, P. and Penandil, P. and Villar, E. and Mitas, N. and Kritharidis, D. and Azcarate, F. and Carballeda, M. The SATURN Approach to SysML-Based HW/SW Codesign. In *VLSI (ISVLSI), 2010 IEEE Computer Society Annual Symposium on*, pages 506–511, IEEE, Los Alamitos, CA, USA, July 2010. 5 pages.

Final Remarks

[cited at p. ix, 5, 35, 36]

- [NASAseh10] National Aeronautics and Space Administration. *Systems Engineering Handbook*. Technical Report, NASA, NASA Headquarters. Whashington, D.C. 20546, December 2007. [cited at p. 9]
- [Nielsen10a] Claus Ballegaard Nielsen. *Dynamic Reconfiguration of Distributed Systems in VDM-RT*. Master's thesis, Aarhus University, December 2010. [cited at p. 30]
- [Nielsen10b] Claus Ballegaard Nielsen. Towards Dynamic Reconfiguration of Distributed Systems in VDM-RT. In *Semantic Issues in VDM: a BCS-FACS and Overture Workshop*, September 2010. [cited at p. 95]
- [Niemann98] Ralf Niemann. *Hardware/software co-design for data flow dominated embedded systems*. Springer, 1998. 244 pages. [cited at p. 13, 91]
- [Ou&04] Jingzhao Ou and Viktor K. Prasanna. Rapid energy estimation of computations on FPGA based soft processors. In *IEEE Intern. SoC Conf, IEEE Computer Society, Washington, DC, USA, 2004*. 4 pages. [cited at p. 5, 35]
- [Ou&05] Jingzhao Ou and Viktor K. Prasanna. MATLAB/Simulink Based Hardware/Software Co-Simulation for Designing Using FPGA Configured Soft Processors. In *Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05) - Workshop 3 - Volume 04*, pages 148.2–, IEEE Computer Society, Washington, DC, USA, 2005. [cited at p. ix, 5, 35]
- [Patel&04] Hire D. Patel and Sandeep K. Dhukla. *SystemC Kernel Extensions for Heterogeneous System Modelling. A Framework for Multi-MoC Modeling and Simulation*. Kluwer Academic Publishers, P.O. Box 17, 3300 AA Dordrecht, The Netherlands, 2004. 172 pages. [cited at p. 10, 24, 26]
- [Perry02] Douglas L. Perry. *VHDL: Programming by Example*. McGraw Hill, fourth edition, 2002. 476 pages. [cited at p. 31]
- [PIC10F320] Microhip Technology Inc. *PIC10(L)F320/322 Data Sheet. 6/8-Pin, High-Performance, Flash Microcontrollers*. Preliminary DS41585A, Microhip Technology Inc., 2355 West Chandler Blvd.Chandler, AZ 85224-6199, 2011. 210 pages. [cited at p. 22]
- [Plat&92] Nico Plat and Hans Toetenel. *A formal transformation from the BSI/VDM-SL concrete syntax to the core abstract syntax*. Technical Report 92-07, Delft University, March 1992. 75 pages. [cited at p. 29]

Chapter 7. Concluding remarks and future work

- [POOSL] SHE/POOSL Website. <http://www.es.ele.tue.nl/she/index.php?select=32>. October 2011. . [cited at p. ix, 28, 29, 37]
- [Prevostini&07] Mauro Prevostini and Elena Zamsa. SysML Profile for SoC Design and SystemC Transformation. May 2007. 7 pages. [cited at p. 5, 35]
- [Ribeiro&11] Augusto Ribeiro and Kenneth Lausdahl and Peter Gorm Larsen. Run-Time Validation of Timing Constraints for VDM-RT Models. In 9th Overture Workshop, June 2011, Limerick, Ireland, 2011. [cited at p. 49]
- [Robert&10] Thomas Robert and Vincent Perrier. CoFluent Methodology for UML. UML SysML MARTE Flow for CoFluent Studio. White Paper, CoFluent Design, 2010. 19 pages. [cited at p. 27]
- [Rosenberg&10] Doug Rosenberg and Sam Mancarella. *Embedded Systems Development using SysML*. Report, Sparx Systems Ltd and ICONIX, 2010. pages. . [cited at p. 25, 26]
- [Schaumont10] Patrick R. Schaumont. *A Practical Introduction to Hardware/Software Codesign*. Springer, 2010. 396 pages. [cited at p. ix, 15, 16, 17, 19, 21]
- [Shaout&09] Adnan Shaout and Ali H. El-Mouse and Khalid Mattar. Specification and Modelling of HW/SW CO-Design for Heterogenous Embedded Systems. July 2009. [cited at p. 1, 12, 13, 15, 19, 24, 91]
- [Simulink] MathWorks. <http://www.mathworks.com/>. October 2011. Simulink official website. [cited at p. x, 5, 27, 259, 260]
- [Speedway11EDK] Introduction to MicroBlaze Hardware Development. Featuring ISE Embedded Design Suite 13. Speedway course material, May 2011. 112 pages. [cited at p. ix, 5, 34]
- [Speedway11SDK] *Developing MicroBlaze Software with Xilinx SDK*. Avnet Inc., August 2011. 113 pages. Speedway course material. [cited at p. 4, 63]
- [Stanton08] Kevin B. Stanton. 802.1AS Tutorial. Presentation, IEEE, November 2008. [cited at p. 81]
- [Sudhakar05] Sudahakar Yalamanchili. *VHDL A starters guide*. Pearson Prentice Hall, Upper Salle Reiver NJ 07458, second edition edition, 2005. 241 pages. [cited at p. 10]

Final Remarks

- [SysML1.2] *Systems Modeling Language (SysML) Specification*. Technical Report Version 1.2, SysML Modelling team, June 2010. 236 pages. <http://www.sysml.org/docs/specs/OMGSysML-v1.2-10-06-02.pdf>. [cited at p. ix, 25]
- [vanderPutten&07] Petrus Henricus Antonius van der Putten and Jeroen Peter Marie Voeten. *Specification of Reactive Hardware/Software Systems. The method Software/Hardware Engineering*. PhD thesis, Technical University of Eindhoven, May 1997. 479 pages. [cited at p. ix, 38]
- [VDMTools] CSK. VDMTools homepage. <http://www.vdmttools.jp/en/>, 2007. pages. . [cited at p. 5, 30]
- [Verhoef05] Marcel Verhoef. *On the Use of VDM++ for Specifying Real-Time Systems*. Proc. First Overture workshop, November 2005. [cited at p. 4, 30]
- [Verhoef08] Marcel Verhoef. *Modeling and Validating Distributed Embedded Real-Time Control Systems*. PhD thesis, Radboud University Nijmegen, 2008. [cited at p. 30] ISBN 978-90-9023705-3
- [Waddington&06] Daniel Waddington and Patrick Lardieri. *Model-Centric Software Development*. *IEEE Computer*, 39(2):28–29, February 2006. 2 pages. . [cited at p. 1, 255]
- [Wolf03] Wolf, Wayne. *A Decade of Hardware/Software Codesign*. 36:38–43, April 2003. [cited at p. 2, 15, 21, 23, 27, 255]
- [XilinxAVB11] Xilinx. *LogiCore IP Ethernet AVB Endpoint v3.1*. DS677, Xilinx, March 2011. pages. [cited at p. 96]
- [XilinxToolFlow11] Xilinx. *Xilinx Tool Flow*. Speedway course material., 2011. 22 pages. [cited at p. 5, 34]

Appendices

Terminology

ASIC

Application Specific Integrated Circuit. Generic integrated circuit that has been tailored for a specific application. A single ASIC might replace several conventional logic integrated circuits, allowing a reduction in both circuit board size and power consumption. [Amos&02].

ASIP

Application Specific Instruction Set Processor. Processors that present a partially reconfigurable Instruction Set Architecture. This allows tuning and optimization of the the processor for certain applications.

AVB

Audio Video Bridging refers to a set of standards developed by the IEEE 802.1 AVB Task Group. It provides the specifications that will allow time-synchronized lo latency streaming services through IEEE 802 networks.

BCAM

Bus Cycle Accurate Model. Model illustrates the communication aspects of the system at the clock cycle level.

BFM

Bus Functional Model. Model that illustrates the communication aspects of the system incorporating no or an approximate notion of time.

CAM

Cycle Accurate Model. Model that illustrates the communication and functional aspects of the system at the clock cycle level.

CCAM

Computational Cycle Accurate Model. Model that illustrates the computational aspects of the system at the clock cycle level.

CLC

Configurable Logic Cell. Reduced programmable silicon area that can implemented functionality at the hardware level defined by the end user. CLC are currently integrated in some commercial microcontrollers enabling the developers the implementation of very simple logical functions on them.

CMOS

Complementary Metal-Oxide Semiconductor. A for of construction of logic monolithic integrated circuits using complementary insulated-gate field-effect transistors in pairs. [Amos&02]

COTS

Commercial Off-The-Shelf. Common term to refer to the commercial availability of a certain product or component.

CSV

Comma Separated Value. Format in which several values are separated by a combination of commas and line breaks.

DUT

Design Under Test. Solution that is being subject of testing activities.

EDA tool

Electronic Design Automation Tool. Software tool used to support the electronic design activities, automating repetitive process involved in the design and speeding up the development process.

Golden Model

Model that is used as a reference during the whole development process

GPIO

General Purpose Input Output. Interface between a processor and external hardware. No specific usage or application is specified, and it can be generally accessed from custom defined logic in order to generate control signals.

IP core

Intellectual Property core. Hardware block implemented by a third party and COTS available. Thoroughly tested and ready for integration in new systems.

ISA

Instruction Set Architecture. Layer between the processor micro-architecture and the software logic. The ISA defines the operations that are available for software execution, as well as other hardware based facilities like available register, memory addressing or interrupt handlers.

Firmcore

IP block implementing functionality via HDL. It has been optimized for a specific target technology. The implementation optimization may be physical layout aware, allowing a highly efficient implementation resulting in improved performance, power or area characteristics. [Cofer&06]. An example of a firmcore block could be a NIOS or Microblaze processors, optimized for Altera and Xilinx platforms respectively.

FPGA

Field Programmable Gate Array is an integrated circuit that can be re-configured by the end user. Its hardware can be implemented by using hardware description languages, like VHDL or Verilog.

Hardcore

IP block implementing functionality at the hardware level. The functionality has been created in fixed-logic at the gate and signal route level rather than within the programmable

FPGA logic fabric. The functionality is fixed at the silicon level during the manufacture of the device. The functionality cannot be removed or modified by the design team. [Cofer&06]. A example of a hardcore block could be an Atmel 8-bit microcontroller.

Hardware/Software Co-design

System Level Design methodology that pursues the derivation of architecture from behaviour and the parallel and cooperative development of hardware and software.

MBSE

Model-Based System Engineering. System engineering approach in which models are guiding the development process, supporting engineering trade-offs evaluation and design decisions.

Partitioning

process/stage of the project in which it is being decided whether a certain functionality should be implemented as a hardware or as a software block.

POOSL

Parallel Object-Oriented Specification Language. Formal language that supports the modelling of concurrent solutions and Object-Oriented techniques. It is used the Software Hardware Engineering methodology.

PSoC

Programmable System-on-Chip. User configurable hardware platform, that allows the incorporation of a number of analog and digital hardware blocks provided by the PSoC manufacturer. Hardware blocks are used by the integrated processor in the PSoC, which can execute a user defined software logic.

PWM

Pulse Width Modulation. Also known as Pulse Duration Modulation. Modulation technique in which the duration of the pulses in a pulse carrier is made to vary in accordance with the instantaneous value of the modulating signal. [Amos&02]

RTC

Real Time Clock. Electronic or electro-mechanical device able to measure physical time.

RTL

Register Transfer Level. Lowest abstraction level used in hardware design. It describes the system by the use of the hardware modelling language VHDL. It enables the developer the performance of very detailed modelling and analysis activities. Due to its low level of abstraction it is one of the most complex representations that can be used in system design.

RTOS

Real Time Operating System. Operating system that offers the possibility of coordinating processes and operations according to physical time limitations.

SAM

System Architectural Model. High level abstraction model of a system. It is normally created at the beginning of the project development activities, and serve as a base model from which more refined TLM models are created.

™SDK

Software Development Kit. Collection of software components that enables the development of additional software.

™SL

Specification Language. Formal or graphical language that allow the representation of system requirements.

™SLD

System Level Design. Engineering approach used in the creation of complex systems. SLD pursues starting the system development at the highest abstraction level possible, as a way to tackle complexity. SLD current approaches are: Component Based Design, Platform Based Design and Hardware/Software Co-Design.

™SM

Specification Model. Model created using a Specification Language, with the purpose of presenting unambiguously the requirements that apply on a certain system.

™Softcore

IP block implementing functionality via HDL with no or minimal optimization for a specific target technology. [Cofer&06]. An example of a softcore could be an ARM core, that can be deployed in several FPGA platforms.

™SysML

System Modelling Language. Graphical modelling language developed by the Object Management Group. It is based on UML and it provides additional constructs for the representation of common problems in the systems engineering domain. SysML is not restricted to hardware or software applications, and it can represents a variety of system components.

™TLM

Transaction Level Model. System model that incorporates no or approximate notion of time. It can detail functional or communication aspects or both. It is an intermediate representation between the SAM models and the CAM models

™UML

Unified Modelling Language. Graphical modelling language developed by the Object Management Group. Widely applied in the software industry, provides constructs to represent object-oriented solutions. Representations can be describing architecture, behaviour and interaction between different software units.

™V model

or V-life cycle. Model of a general development process, that shows the existing relationship between analysis, design and implementation and validation and verification phases.

™VDM

Vienna Development Method. Formal method enabling the modelling of systems. It can be used along all the phases described by the V model. The VDM method is currently supporting three languages: VDM-SL, VDM++ and VDM-RT. These languages cover the specification, object oriented design and real-time modelling respectively.

™VHDL

Very-high-speed Hardware Description Language. Low level language used for the creation

of hardware components. The implementation of these hardware components can be carried out in different platforms (FPGAs, ASICs, ...).

Servo case study VDM models

This appendix contains the VDM models for the servo case study. Section B.1 presents the sequential VDM models. Section B.2 presents the concurrent VDM models. Finally, sections B.12, B.13, B.14 present the produced VDM-RT models.

B.1. Sequential VDM model

B.1.1 Clock

```
1
2 class Clock
3
4 instance variables
5
6   private tickNum : nat := 0;
7
8 operations
9
10 public tick : () ==> ()
11 tick() ==
12 (
13   tickNum := tickNum +1;
14 );
15
16 public getSimTime: () ==> nat
17 getSimTime() ==
18   return tickNum;
19
20
21 end Clock
```

B.1.2 Controller

```

1
2 class Controller
3
4 instance variables
5
6   private pwmGenerator : PWMgenerator;
7   private outputInterface : gpioInterface;
8
9 operations
10
11 public Controller: gpioInterface * PWMgenerator ==> Controller
12 Controller (gpio,p) ==
13 (
14   outputInterface := gpio;
15   -- Create a position logger
16   pwmGenerator := p;
17 );
18
19 public getOutputInterface: () ==> gpioInterface
20 getOutputInterface() ==
21   return outputInterface;
22
23 -- Operations running in the controller
24 public timeStep: nat ==> ()
25 timeStep(time) ==
26 (
27   skip;
28 );
29
30 end Controller

```

B.1.3 Environment

```

1
2 class Environment
3
4 instance variables
5
6   clk : Clock;
7   simTimeMax : nat := 1200;
8   currentSimTime : nat := 0;
9
10  outputInterface : gpioInterface;

```

Environment

```
11 controller : Controller;
12
13 sampler1 : Sampler;
14 posLog : PositionLogger;
15
16 pwmUnit : PWMgenerator;
17
18 operations
19
20 public run : () ==> ()
21 run() ==
22 (
23   -- ### Simulation setup
24   -- Create a clock
25   clk := new Clock();
26
27   -- Define a microcontroller
28
29   posLog := new PositionLogger();
30
31   -- Attach the gpio generating the PWM signal to the environment
32   outputInterface := new gpioInterface(2);
33   -- Create a PWMgenerator and attach it to the output
34   pwmUnit := new PWMgenerator(outputInterface,1,posLog);
35
36   controller := new Controller(outputInterface,pwmUnit);
37   -- Set the initial pulse duration to 10 ms
38   pwmUnit.pulseDuration(10);
39
40   -- Create a virtual scope and attach it to the generated
41   -- PWM output.
42   -- Using channel 1 from the GPIO block.
43   sampler1 := new Sampler(outputInterface,1);
44
45   -- ### Run simulation
46   while (clk.getSimTime() <= simTimeMax) do
47   (
48     currentSimTime := clk.getSimTime();
49     IO`print("\nSimulation time step: ");
50     IO`print(clk.getSimTime());
51
52     controller.timeStep(currentSimTime);
53     pwmUnit.timeStep();
54
55
56     sampler1.sample();
57     clk.tick();
58   );
59
```

Appendix B. Servo case study VDM models

```
60     -- ### Simulation tear down
61     sampler1.dumpToFile();
62     pwmUnit.recoverLogs();
63     sampler1.dumpClock();
64
65     IO`print("## Model Over ##");
66 );
67
68 end Environment
```

B.1.4 PWMgenerator

```
1
2 class PWMgenerator
3
4 values
5
6     -- Number of clock ticks per cycle
7     private ticksPerPulse : nat = 200;
8
9     -- 10 ticks = 1 ms
10    private msFactor : nat = 10;
11
12    -- Signal frequency in Hz
13    private frequency : nat = 50;
14
15 instance variables
16
17    -- Used GPIO block and channel number
18    private output : gpioInterface;
19    private channel : nat := 0;
20
21    -- Internal counter used to define the
22    -- transition from high to low inside
23    -- each cycle.
24
25    private tickCounter : nat := 0;
26    inv tickCounter >= 0 and tickCounter <= ticksPerPulse;
27    -- tickCounter --> is set ot 0 at the beginning of A
28    --             --> incremented each clock tick
29    --             --> when tickCounter reaches its limit
30    --             the line is asserted for the rest of the
31    --             cycle.
32    -- The end of A and beggining of B depends on the desired
33    -- PWM signal
34    --
```

PWMgenerator

```

35  -- A      B      A      B
36  -- <--><-----><--><----->
37  --      |      |      |      |
38  --      |      |      |      |
39  --      |      |      |      |
40
41  private tickLimit : nat := 0;
42  inv tickLimit > 0 and tickLimit < ticksPerPulse;
43
44  -- Higher duty cycle values might damage the servo
45  inv getDutyCycle() < 10;
46
47  public posLog : PositionLogger;
48
49  operations
50
51  -- Constructor
52  -- Associates a GPIO block and a channel.
53  public PWMgenerator : gpioInterface *
54                      nat *
55                      PositionLogger ==>
56                      PWMgenerator
57  PWMgenerator(out, ch, posL) ==
58  (
59  output := out;
60  channel := ch;
61  tickLimit := ticksPerPulse/2;
62  posLog := posL;
63  )
64  post output.getWidth() >= channel and
65  not channel = 0 and
66  getDutyCycle() = 0.5;
67
68  public calculatePosition: () ==> real
69  calculatePosition() ==
70  return (180 * ((tickLimit/msFactor) -1));
71
72
73  -- Takes in account the duty cycle and generates the signal
74  private generateCycle: () ==> ()
75  generateCycle() ==
76  cases tickCounter:
77  -- Start of the new cycle
78  (0) -> setOn(),
79  -- Time limit for the high level has been reached
80  (tickLimit) -> setOff(),
81  -- Keep the line as it is
82  others -> skip
83  end;

```

Appendix B. Servo case study VDM models

```
84 |
85 |
86 | public timeStep: () ==> ()
87 | timeStep() ==
88 | (
89 |   -- Generate the part of the continuous signal
90 |   -- that corresponds to this time step.
91 |   generateCycle();
92 |   IO`print("\n## Calculated position in degrees: ");
93 |   IO`print(calculatePosition());
94 |   posLog.pushValue(calculatePosition());
95 |
96 |   -- Update the pulse counter in the case the resolution
97 |   -- limit is reached
98 |   if tickCounter >= ticksPerPulse then
99 |     tickCounter := 0
100 |   else
101 |     tickCounter := tickCounter +1;
102 | );
103 |
104 | -- Sets pulse duration in milliseconds
105 | public pulseDuration : nat ==> ()
106 | pulseDuration(d) ==
107 |   tickLimit := d;
108 |
109 | -- Assert high the associated GPIO channel
110 | private setOn: () ==> ()
111 | setOn() ==
112 |   output.setHigh(channel);
113 |
114 | -- Assert low the associated GPIO channel
115 | private setOff: () ==> ()
116 | setOff() ==
117 |   output.setLow(channel);
118 |
119 | public gettickLimit: () ==> nat
120 | gettickLimit() ==
121 |   return tickLimit;
122 |
123 | public setLimit : nat ==> ()
124 | setLimit(l) ==
125 |   tickLimit := l;
126 |
127 | public recoverLogs: () ==> ()
128 | recoverLogs() ==
129 | (
130 |   posLog.dumpToFile();
131 | );
132 |
```


PositionLogger

```
133 functions
134
135 -- Calculates the current duty cycle of the signal
136 -- defined by timeHigh/timeCycle
137 public getDutyCycle : () -> real
138 getDutyCycle() ==
139   gettickLimit()/ticksPerPulse
140 pre ticksPerPulse > 0;
141
142 end PWMgenerator
```

B.1.5 PositionLogger

```
1
2 class PositionLogger
3
4 types
5
6   private OutputTP = int * int;
7
8 instance variables
9
10  private numValues : nat := 0;
11  private sampledValues : seq of real := [];
12
13  private writeResult : bool := false;
14  private io : IO := new IO();
15
16 operations
17
18 public pushValue: real ==> ()
19 pushValue(angle) ==
20 (
21   sampledValues:= sampledValues ^ [angle];
22 );
23
24 public dumpToFile:() ==> ()
25 dumpToFile () ==
26 (
27   IO`print(sampledValues);
28
29   for all r in set {1,..., len(sampledValues)}
30   do
31   (
32     writeResult := io.fwriteval[real]("position.csv",
33                                     sampledValues(r),
```

Appendix B. Servo case study VDM models

```
34         <append>);
35     writeResult := io.fecho("position.csv", ",1", <append>);
36     writeResult := io.fecho("position.csv", "\n", <append>);
37 );
38
39 );
40
41 end PositionLogger
```

B.1.6 Sampler

```
1
2 class Sampler
3
4 types
5
6     OutputTP = int * int;
7
8 instance variables
9
10    gpioUnit : gpioInterface;
11    sampledValues : seq of bool := [];
12    bitToSample : nat;
13    numValues : nat := 0;
14    inv numValues = len sampledValues;
15
16    toFile : seq of nat := [];
17
18    writeResult : bool := false;
19    io : IO := new IO();
20
21
22 operations
23
24 public sample : () ==> ()
25 sample () ==
26     atomic
27     (
28         sampledValues := sampledValues ^
29             [gpioUnit.getState(bitToSample)];
30         numValues := numValues + 1;
31     );
32
33 public Sampler : gpioInterface * nat ==> Sampler
34 Sampler(gpioBlock,output) ==
35 (
```

Sampler

```
36 | gpioUnit := gpioBlock;
37 | bitToSample := output;
38 | );
39 |
40 | public dumpToFile : () ==> ()
41 | dumpToFile () ==
42 | (
43 |   IO`print(sampledValues);
44 |
45 |   for all r in set {1,...,numValues}
46 |   do
47 |   (
48 |     IO`print(r);
49 |     writeResult := io.fwriteval[int]("signal.csv", r, <append>);
50 |     if sampledValues(r) then
51 |     (
52 |       --IO`print(",1");
53 |       writeResult := io.fecho("signal.csv", ",1", <append>)
54 |     )
55 |     else
56 |     (
57 |       --IO`print(",0");
58 |       writeResult := io.fecho("signal.csv",",0", <append>)
59 |     );
60 |
61 |     writeResult := io.fecho("signal.csv","\n", <append>);
62 |     --IO`print("\n");
63 |   );
64 | );
65 |
66 | public dumpClock : () ==> ()
67 | dumpClock () ==
68 | (
69 |
70 |   for all r in set {1,...,numValues}
71 |   do
72 |   (
73 |     writeResult := io.fwriteval[int]("clock.csv", r, <append>);
74 |     writeResult := io.fecho("clock.csv", ",0", <append>);
75 |     writeResult := io.fecho("clock.csv","\n", <append>);
76 |
77 |     writeResult := io.fwriteval[int]("clock.csv", r, <append>);
78 |     writeResult := io.fecho("clock.csv", ",1", <append>);
79 |     writeResult := io.fecho("clock.csv","\n", <append>);
80 |
81 |     writeResult := io.fwriteval[real]("clock.csv",
82 |                                     r+0.5,
83 |                                     <append>);
84 |     writeResult := io.fecho("clock.csv", ",1", <append>);
```

Appendix B. Servo case study VDM models

```
85     writeResult := io.fecho("clock.csv","\n", <append>);
86
87     writeResult := io.fwriteval[real]("clock.csv",
88                                     r+0.5,
89                                     <append>);
90     writeResult := io.fecho("clock.csv", ",0", <append>);
91     writeResult := io.fecho("clock.csv","\n", <append>);
92 )
93 );
94
95 public assignOutput : gpioInterface ==> ()
96 assignOutput (gpioU) ==
97 (
98     gpioUnit := gpioU;
99 );
100
101 end Sampler
```

B.1.7 Timer

```
1 class Timer
2
3 instance variables
4
5     freq : nat;
6     flag : bool := true;
7
8 operations
9
10 public Timer: (nat) ==> Timer
11 Timer(fq) ==
12 (
13     freq := fq;
14 );
15
16 private wait: () ==> ()
17 wait() ==
18 (
19 );
20 );
21
22 private toggleFlag: () ==> ()
23 toggleFlag() ==
24     flag := not flag;
25
26 public getFlag: () ==> bool
```

gpioInterface

```
27 | getFlag() ==
28 |     return flag;
29 |
30 | public activeBehaviour:() ==> ()
31 | activeBehaviour() ==
32 | (
33 |
34 | );
35 |
36 | thread
37 |
38 | while true do
39 | (
40 |     wait();
41 |     toggleFlag();
42 |     IO`print("\nFlag state: ");
43 |     IO`print(flag);
44 | );
45 |
46 | sync
47 |     mutex(toggleFlag,getFlag);
48 |
49 | end Timer
```

B.1.8 gpioInterface

```
1 |
2 | class gpioInterface
3 |
4 | instance variables
5 |
6 |     width : nat := 0;
7 |     state : seq of bool := [];
8 |     inv len state = width;
9 |
10 | operations
11 |
12 | public gpioInterface: nat ==> gpioInterface
13 |     gpioInterface(w) ==
14 | (
15 |     width := w;
16 |
17 |     for all r in set {1,...,width}
18 |         do state := state ^ [false];
19 |
20 |     IO`print(state);
```

```

21 |
22 | )
23 | post width = len state and
24 |   forall i in set {1,...,len state} & state(i) = false;
25 |
26 |
27 | public setHigh: nat ==> ()
28 | setHigh(i) ==
29 |   state(i) := true
30 | pre i <= len state;
31 |
32 | public setLow: nat ==> ()
33 | setLow(i) ==
34 |   state(i) := false
35 | pre i <= len state;
36 |
37 | public toggleBit: nat ==> ()
38 | toggleBit(i) ==
39 |   if state(i)
40 |   then state(i) := not state(i)
41 |   else state(i) := true
42 | pre i <= len state;
43 |
44 | public getState: nat ==> bool
45 | getState(i) ==
46 |   return state(i)
47 | pre i <= len state;
48 |
49 | public showState: () ==> ()
50 | showState() ==
51 |   for all r in set {1,...,width}
52 |   do showSingle(r);
53 |
54 |
55 | public getWidth: () ==> nat
56 | getWidth() ==
57 |   return width;
58 |
59 | public showSingle: nat ==> ()
60 | showSingle(i) ==
61 | (
62 |   IO`print("\n bit ");
63 |   IO`print(i);
64 |   IO`print(": ");
65 |   IO`print(getState(i));
66 | );
67 |
68 | sync
69 |   mutex(setLow, setHigh);

```

```

70  mutex(setHigh,getState);
71  mutex(setLow,getState);
72
73  end gpioInterface

```

B.2. Concurrent VDM model

B.3. Clock

```

1
2  class Clock
3
4  instance variables
5
6  private tickNum : nat := 0;
7
8  operations
9
10 public tick : () ==> ()
11 tick() ==
12 (
13   tickNum := tickNum +1;
14 );
15
16 public getSimTime: () ==> nat
17 getSimTime() ==
18   return tickNum;
19
20 thread
21 while true do
22 (
23   tick();
24   Environment `timerRef.WaitRelative(1);
25 );
26
27 end Clock

```

B.4. Controller

```

1

```

```

2 class Controller
3
4 instance variables
5
6   private pwmGenerator : PWMgenerator;
7   private outputInterface : gpioInterface;
8
9 operations
10
11 public Controller: gpioInterface * PWMgenerator ==> Controller
12 Controller (gpio,p) ==
13 (
14   outputInterface := gpio;
15   -- Create a position logger
16   pwmGenerator := p;
17 );
18
19 public getOutputInterface: () ==> gpioInterface
20 getOutputInterface() ==
21   return outputInterface;
22
23 -- Operations running in the controller
24 public timeStep: () ==> ()
25 timeStep() ==
26 (
27   skip;
28 );
29
30 thread
31
32 while true do
33 ( timeStep();
34   Environment `timerRef.WaitRelative(1);
35 );
36
37 end Controller

```

B.5. Environment

```

1
2 class Environment
3
4 instance variables
5

```


Environment

```
6   clk : Clock;
7   simTimeMax : nat := 2400;
8   currentSimTime : nat := 0;
9
10  outputInterface : gpioInterface;
11  controller : Controller;
12
13  public sampler1 : Sampler;
14  posLog : PositionLogger;
15
16  pwmUnit : PWMgenerator;
17
18  public static timerRef : TimeStamp := new TimeStamp(4);
19
20  operations
21
22  public run : () ==> ()
23  run() ==
24  (
25    -- ### Simulation setup
26    -- Create a clock
27    clk := new Clock();
28
29    -- Define a microcontroller
30
31    posLog := new PositionLogger();
32
33    -- Attach the gpio generating the PWM signal to the environment
34    outputInterface := new gpioInterface(2);
35    -- Create a PWMgenerator and attach it to the output
36    pwmUnit := new PWMgenerator(outputInterface,1,posLog);
37
38    controller := new Controller(outputInterface,pwmUnit);
39    -- Set the initial pulse duration to 10 ms
40    pwmUnit.pulseDuration(10);
41
42    -- Create a virtual scope and attach it to the generated
43    -- PWM output.
44    -- Using channel 1 from the GPIO block.
45    sampler1 := new Sampler(outputInterface,1);
46
47    start(clk);
48    start(pwmUnit);
49    start(controller);
50    start(sampler1);
51
52    -- ### Run simulation
53    while (timerRef.GetTime() <= simTimeMax) do
54    (
```

Appendix B. Servo case study VDM models

```
55     timerRef.WaitRelative(1);
56     wait();
57 );
58
59     -- ### Simulation tear down
60     sampler1.dumpToFile();
61     pwmUnit.recoverLogs();
62     sampler1.dumpClock();
63
64     IO`print("## Model Over ##");
65 );
66
67
68 public wait: () ==> ()
69 wait() ==
70     skip;
71
72 end Environment
```

B.6. PWMgenerator

```
1
2 class PWMgenerator
3
4 values
5
6     -- Number of clock ticks per cycle
7     private ticksPerPulse : nat = 200;
8
9     -- 10 ticks = 1 ms
10    private msFactor : nat = 10;
11
12    -- Signal frequency in Hz
13    private frequency : nat = 50;
14
15 instance variables
16
17    -- Used GPIO block and channel number
18    private output : gpioInterface;
19    private channel : nat := 0;
20
21    -- Internal counter used to define the
22    -- transition from high to low inside
23    -- each cycle.
```

PWMgenerator

```

24
25 private tickCounter : nat := 0;
26 inv tickCounter >= 0 and tickCounter <= ticksPerPulse;
27 -- tickCounter --> is set ot 0 at the beginning of A
28 -- --> incremented each clock tick
29 -- --> when tickCounter reaches its limit
30 -- the line is asserted for the rest of the
31 -- cycle.
32 -- The end of A and beggining of B deppends on the desired
33 -- PWM signal
34 --
35 -- A      B      A      B
36 -- <--><-----><--><----->
37 --  _      _
38 -- |  |    |  |
39 -- |  |_____|  |  |_____
40
41 private tickLimit : nat := 0;
42 inv tickLimit > 0 and tickLimit < ticksPerPulse;
43
44 -- Higher duty cycle values might damage the servo
45 inv getDutyCycle() < 10;
46
47 public posLog : PositionLogger;
48
49 operations
50
51 -- Constructor
52 -- Associates a GPIO block and a channel.
53 public PWMgenerator : gpioInterface *
54     nat *
55     PositionLogger ==>
56     PWMgenerator
57 PWMgenerator(out, ch, posL) ==
58 (
59   output := out;
60   channel := ch;
61   tickLimit := ticksPerPulse/2;
62   posLog := posL;
63 )
64 post output.getWidth() >= channel and
65     not channel = 0 and
66     getDutyCycle() = 0.5;
67
68 public calculatePosition: () ==> real
69 calculatePosition() ==
70     return (180 * ((tickLimit/msFactor) -1));
71
72

```

Appendix B. Servo case study VDM models

```
73 -- Takes in account the duty cycle and generates the signal
74 private generateCycle: () ==> ()
75 generateCycle() ==
76   cases tickCounter:
77     (0) -> setOn(),
78     (tickLimit) -> setOff(),
79     others -> skip
80   end;
81
82
83 public timeStep: () ==> ()
84 timeStep() ==
85 (
86   -- Generate the part of the continuous signal
87   -- that corresponds to this time step.
88   generateCycle();
89   IO`print("\n## Calculated position in degrees: ");
90   IO`print(calculatePosition());
91   posLog.pushValue(calculatePosition());
92
93   -- Update the pulse counter in the case the resolution
94   -- limit is reached
95   if tickCounter >= ticksPerPulse then
96     tickCounter := 0
97   else
98     tickCounter := tickCounter + 1;
99
100   Environment`timerRef.WaitRelative(1);
101 );
102
103 -- Sets pulse duration in milliseconds
104 public pulseDuration : nat ==> ()
105 pulseDuration(d) ==
106   tickLimit := d;
107
108 -- Assert high the associated GPIO channel
109 private setOn: () ==> ()
110 setOn() ==
111   output.setHigh(channel);
112
113 -- Assert low the associated GPIO channel
114 private setOff: () ==> ()
115 setOff() ==
116   output.setLow(channel);
117
118 public gettickLimit: () ==> nat
119 gettickLimit() ==
120   return tickLimit;
121
```

PositionLogger

```
122 public setLimit : nat ==> ()
123 setLimit(1) ==
124   tickLimit := 1;
125
126 public recoverLogs: () ==> ()
127 recoverLogs() ==
128 (
129   posLog.dumpToFile();
130 );
131
132 functions
133
134 -- Calculates the current duty cycle of the signal
135 -- defined by timeHigh/timeCycle
136 public getDutyCycle : () -> real
137 getDutyCycle() ==
138   gettickLimit()/ticksPerPulse
139 pre ticksPerPulse > 0;
140
141 thread
142
143 while true do
144 (
145   timeStep();
146   Environment `timerRef.WaitRelative(1);
147 );
148
149
150 end PWMgenerator
```

B.7. PositionLogger

```
1
2 class PositionLogger
3
4 types
5
6   private OutputTP = int * int;
7
8 instance variables
9
10  private numValues : nat := 0;
11  private sampledValues : seq of real := [];
12
```

```

13  private writeResult : bool := false;
14  private io : IO := new IO();
15
16  operations
17
18  public pushValue: real ==> ()
19  pushValue(angle) ==
20  (
21    sampledValues:= sampledValues ^ [angle];
22  );
23
24  public dumpToFile:() ==> ()
25  dumpToFile () ==
26  (
27    IO`print(sampledValues);
28
29    for all r in set {1,..., len(sampledValues)}
30    do
31    (
32      writeResult := io.fwriteval[real]("position.csv",
33                                     sampledValues(r),
34                                     <append>);
35      writeResult := io.fecho("position.csv", ",1", <append>);
36      writeResult := io.fecho("position.csv","\n", <append>);
37    );
38  );
39
40
41  end PositionLogger

```

B.8. Sampler

```

1
2  class Sampler
3
4  types
5
6  OutputTP = int * int;
7
8  instance variables
9
10  gpioUnit : gpioInterface;
11  sampledValues : seq of bool := [];
12  bitToSample : nat;

```

Sampler

```
13 numValues : nat := 0;
14 -- inv numValues = len sampledValues;
15
16 toFile : seq of nat := [];
17
18 writeResult : bool := false;
19 io : IO := new IO();
20
21
22 operations
23
24 public sample : () ==> ()
25 sample () ==
26 (
27   atomic
28   (
29     sampledValues := sampledValues ^
30                       [gpioUnit.getState(bitToSample)];
31     numValues := numValues +1;
32   );
33 );
34
35 public Sampler : gpioInterface * nat ==> Sampler
36 Sampler(gpioBlock,output) ==
37 (
38   gpioUnit := gpioBlock;
39   bitToSample := output;
40 );
41
42 public dumpToFile : () ==> ()
43 dumpToFile () ==
44 (
45   IO`print(sampledValues);
46
47   for all r in set {1,...,numValues}
48   do
49   (
50     IO`print(r);
51     writeResult := io.fwriteval[int]("signal.csv", r, <append>);
52     if sampledValues(r) then
53     (
54       --IO`print(",1");
55       writeResult := io.fecho("signal.csv", ",1", <append>)
56     )
57     else
58     (
59       --IO`print(",0");
60       writeResult := io.fecho("signal.csv",",0", <append>)
61     );
```

Appendix B. Servo case study VDM models

```
62
63     writeResult := io.fecho("signal.csv","\n", <append>);
64     --IO `print ("\n");
65     );
66 );
67
68 public dumpClock : () ==> ()
69 dumpClock () ==
70 (
71
72     for all r in set {1,...,numValues}
73     do
74     (
75         writeResult := io.fwriteval[int]("clock.csv", r, <append>);
76         writeResult := io.fecho("clock.csv", ",0", <append>);
77         writeResult := io.fecho("clock.csv","\n", <append>);
78
79         writeResult := io.fwriteval[int]("clock.csv", r, <append>);
80         writeResult := io.fecho("clock.csv", ",1", <append>);
81         writeResult := io.fecho("clock.csv","\n", <append>);
82
83         writeResult := io.fwriteval[real]("clock.csv",
84                                     r+0.5,
85                                     <append>);
86         writeResult := io.fecho("clock.csv", ",1", <append>);
87         writeResult := io.fecho("clock.csv","\n", <append>);
88
89         writeResult := io.fwriteval[real]("clock.csv",
90                                     r+0.5,
91                                     <append>);
92         writeResult := io.fecho("clock.csv", ",0", <append>);
93         writeResult := io.fecho("clock.csv","\n", <append>);
94     )
95 );
96
97 public assignOutput : gpioInterface ==> ()
98 assignOutput (gpioU) ==
99 (
100     gpioUnit := gpioU;
101 );
102
103 thread
104     while true do
105     (
106         sample();
107         Environment `timerRef.WaitRelative(1);
108     );
109
110 end Sampler
```


B.9. TimeStamp

```

1 class TimeStamp
2
3 values
4
5 public stepLength : nat = 1;
6
7 instance variables
8
9 currentTime : nat := 0;
10 wakeupMap   : map nat to [nat] := {|->};
11 barrierCount : nat1;
12
13 operations
14
15 public TimeStamp : nat1 ==> TimeStamp
16 TimeStamp(count) ==
17   barrierCount := count;
18
19 public WaitRelative : nat ==> ()
20 WaitRelative(val) ==
21   WaitAbsolute(currentTime + val);
22
23 public WaitAbsolute : nat ==> ()
24 WaitAbsolute(val) == (
25   AddToWakeupMap(threadid, val);
26   -- Last to enter the barrier notifies the rest.
27   BarrierReached();
28   -- Wait till time is up
29   Awake();
30 );
31
32 BarrierReached : () ==> ()
33 BarrierReached() ==
34 (
35   while (card dom wakeupMap = barrierCount) do
36     (
37       currentTime := currentTime + stepLength;
38       let threadSet : set of nat = {th | th in set dom wakeupMap
39         & wakeupMap(th) <> nil and
40         wakeupMap(th) <= currentTime }
41   in

```

Appendix B. Servo case study VDM models

```
42   for all t in set threadSet
43   do
44     wakeUpMap := {t} <-: wakeUpMap;
45   );
46 )
47 post forall x in set rng wakeUpMap & x = nil or x >= currentTime;
48
49 AddToWakeUpMap : nat * [nat] ==> ()
50 AddToWakeUpMap(tId, val) ==
51   wakeUpMap := wakeUpMap ++ { tId |-> val };
52
53 public NotifyThread : nat ==> ()
54 NotifyThread(tId) ==
55   wakeUpMap := {tId} <-: wakeUpMap;
56
57 public GetTime : () ==> nat
58 GetTime() ==
59   return currentTime;
60
61 Awake: () ==> ()
62 Awake() == skip;
63
64 public ThreadDone : () ==> ()
65 ThreadDone() ==
66   AddToWakeUpMap(threadid, nil);
67
68 sync
69   per Awake => threadid not in set dom wakeUpMap;
70
71   mutex(AddToWakeUpMap);
72   mutex(NotifyThread);
73   mutex(BarrierReached);
74
75   mutex(AddToWakeUpMap, NotifyThread);
76   mutex(AddToWakeUpMap, BarrierReached);
77   mutex(NotifyThread, BarrierReached);
78
79   mutex(AddToWakeUpMap, NotifyThread, BarrierReached);
80
81 end TimeStamp
```

B.10. Timer

```
1 class Timer
```

```
2
3 instance variables
4
5   freq : nat;
6   flag : bool := true;
7
8 operations
9
10 public Timer: (nat) ==> Timer
11 Timer(fq) ==
12 (
13   freq := fq;
14 );
15
16 private wait: () ==> ()
17 wait() ==
18 (
19
20 );
21
22 private toggleFlag: () ==> ()
23 toggleFlag() ==
24   flag := not flag;
25
26 public getFlag: () ==> bool
27 getFlag() ==
28   return flag;
29
30 public activeBehaviour:() ==> ()
31 activeBehaviour() ==
32 (
33
34 );
35
36 thread
37
38 while true do
39 (
40   wait();
41   toggleFlag();
42   IO`print("\nFlag state: ");
43   IO`print(flag);
44 );
45
46 sync
47   mutex(toggleFlag,getFlag);
48
49 end Timer
```

B.11. gpioInterface

```

1
2 class gpioInterface
3
4 instance variables
5
6   width : nat := 0;
7   state : seq of bool := [];
8   inv len state = width;
9
10 operations
11
12 public gpioInterface: nat ==> gpioInterface
13   gpioInterface(w) ==
14   (
15     width := w;
16
17     for all r in set {1,...,width}
18     do state := state ^ [false];
19
20     IO`print(state);
21
22   )
23 post width = len state and
24   forall i in set {1,...,len state} & state(i) = false;
25
26
27 public setHigh: nat ==> ()
28   setHigh(i) ==
29     state(i) := true
30 pre i <= len state;
31
32 public setLow: nat ==> ()
33   setLow(i) ==
34     state(i) := false
35 pre i <= len state;
36
37 public toggleBit: nat ==> ()
38   toggleBit(i) ==
39     if state(i)
40     then state(i) := not state(i)
41     else state(i) := true
42 pre i <= len state;
43
44 public getState: nat ==> bool
45   getState(i) ==
46     return state(i)

```

```

47 pre i <= len state;
48
49 public showState: () ==> ()
50 showState() ==
51   for all r in set {1,...,width}
52   do showSingle(r);
53
54
55 public getWidth: () ==> nat
56 getWidth() ==
57   return width;
58
59 public showSingle: nat ==> ()
60 showSingle(i) ==
61 (
62   IO`print("\n bit ");
63   IO`print(i);
64   IO`print(": ");
65   IO`print(getState(i));
66 );
67
68 sync
69   mutex(setLow, setHigh);
70   mutex(setHigh, getState);
71   mutex(setLow, getState);
72
73 end gpioInterface

```

B.12. Real-time VDM model: Software based solution

B.12.1 CPUloader

```

1 class Controller
2
3 -- Definition of the CPU loader running in a controller
4
5 values
6   -- Execution limit for the CPU loader
7   executionLimit : nat = 90;
8
9 instance variables
10
11   -- Switch time mark. Initialize to a higher value
12   -- in the case a discrete control is used from the

```

Appendix B. Servo case study VDM models

```

13  -- loader.
14  switchTimeMark :nat := 0;
15
16  progressCounter : nat := 0;
17  pwmref: PWMunit;
18  cSent :bool := false;
19
20  operations
21
22  -- Specify the switch time mark in the case several
23  -- control comands are issued to the controller
24  public setSwitchTimeMark : nat ==> ()
25  setSwitchTimeMark(tm) ==
26    switchTimeMark := tm;
27
28  -- Time consumer operation loading the CPU. The
29  -- function is not performing any computation, just
30  -- taking CPU time.
31  public timeConsumer : () ==> ()
32  timeConsumer () ==
33  (
34    -- Simulated operations performed during
35    -- the first 9 iterations.
36    cases progressCounter:
37      (1) -> duration(1E6) timeEater(),
38      (2) -> duration(5E6) timeEater(),
39      (3) -> duration(10E6) timeEater(),
40      (4) -> duration(15E6) timeEater(),
41      (5) -> duration(20E6) timeEater(),
42      (6) -> duration(25E6) timeEater(),
43      (7) -> duration(30E6) timeEater(),
44      (8) -> duration(40E6) timeEater(),
45      (9) -> duration(50E6) timeEater(),
46      others -> ()
47    end;
48
49    -- Increment the progress counter
50    progressCounter := progressCounter + 1;
51
52    -- Send a control signal to the hardware block
53    -- in the case a switchTimeMark has been specified
54    if time > switchTimeMark and not cSent then
55      (
56        controlSignal();
57        IO`print("#] PWM generator notified at time: ");
58        IO`print(time); IO`print(" [#]\n");
59        cSent := true;
60      );
61 );

```

Deployment

```
62
63 -- Issues a change in the continuous time hardware
64 -- signal generation.
65 public controlSignal : () ==> ()
66 controlSignal () ==
67 (
68   pwmref.setPeriod(2);
69 );
70
71 -- Establish the reference to the PWM unit
72 public setPWM : PWMunit ==> ()
73 setPWM(p) ==
74   pwmref := p;
75
76 -- Do nothing
77 public timeEater : () ==> ()
78 timeEater() ==
79   skip;
80
81 thread
82   -- Procedural thread running unit completion
83   while progressCounter < executionLimit do timeConsumer();
84
85 sync
86   -- Ensures that only one timeConsumer instance is running
87   -- at a certain point of time.
88   mutex(timeConsumer);
89
90 end Controller
```

B.12.2 Deployment

```
1 system Deployment
2
3 instance variables
4
5   -- Definition of a controller unit, acting as CPU
6   controller : CPU := new CPU(<FCFS>, 1E9);
7
8   -- Definition of a hardware GPIO block
9   gpioBlock : CPU := new CPU(<FCFS>, 100E6);
10
11   -- Associate the controller to the GPIO block
12   controlRegister : BUS := new BUS(<CSMACD>,
13                                   72E13,
14                                   {controller,gpioBlock});
```

Appendix B. Servo case study VDM models

```
15
16  -- Static definition of the deployable objects
17  public static loader : Controller := new Controller();
18  public static pwm : PWMunit := new PWMunit();
19  public static gpio : gpioInterface := new gpioInterface(2);
20
21  operations
22
23  -- Setup and actual deployment
24  public Deployment : () ==> Deployment
25  Deployment () ==
26  (
27    -- Set a reference to the pwm generator from the loader
28    loader.setPWM(pwm);
29    -- Set a reference to the GPIO block from the PWMgenerator
30    pwm.setGPIO(gpio);
31
32    -- Deploy the active objects in different CPUs
33    controller.deploy(loader);
34    controller.deploy(pwm);
35    gpioBlock.deploy(gpio);
36
37  );
38
39  end Deployment
```

B.12.3 Environment

```
1  class Environment
2
3  values
4
5    -- Time limit mark in the case two control signals
6    -- are going to be generated.
7    private timeLimitMark : nat = 0;
8
9  instance variables
10
11    -- Create a sampler
12    public static sampler1 : Sampler := new Sampler();
13
14  operations
15
16  -- Run operation, simulation entry point
17  public run : () ==> ()
18  run() ==
```


PWMgenerator

```
19 | (  
20 |   -- Associate the limit time mark for the loader  
21 |   Deployment `loader.setSwitchTimeMark(timeLimitMark);  
22 |  
23 |   -- Associate the sampler to the PWM generator  
24 |   Deployment `pwm.outLog := sampler1;  
25 |  
26 |   -- Start the threads  
27 |   start (Deployment `pwm);  
28 |   start (Deployment `loader);  
29 |  
30 |   wait(); -- Block until threads are done  
31 |  
32 |   -- Dump the sampled signals to a file  
33 |   sampler1.dumpSignal();  
34 |  
35 |   -- Simulation tear down  
36 |   printAndBye();  
37 |  
38 | );  
39 |  
40 | -- Simulation tear down  
41 | public printAndBye: () ==> ()  
42 | printAndBye() ==  
43 | (  
44 |   IO `print ("[#] Generated signal: ");  
45 |   IO `print (Deployment `pwm.outLog.log);  
46 |   IO `print (" [#]");  
47 |   IO `print ("\n[1] Model over");  
48 |   IO `print ("\n[2] Logs generated");  
49 | );  
50 |  
51 | -- Wait until it is down  
52 | public wait : () ==> ()  
53 | wait() ==  
54 |   skip;  
55 |  
56 | sync  
57 |   -- Sync predicate for the wait operation  
58 |   per wait => Deployment `pwm.getFlag();  
59 |  
60 | end Environment
```

B.12.4 PWMgenerator

```
1 | class PWMunit
```

Appendix B. Servo case study VDM models

```
2
3 instance variables
4
5   -- References to the sampler and the
6   -- controller
7   public outLog : Sampler;
8   public controller : Controller;
9
10  -- Allows the generation of two different
11  -- predefined signals
12  private periodSelector : nat := 1;
13
14  -- Sets a reference to the GPIO interface
15  public outputInt : gpioInterface;
16
17  -- Done flag updated when operation completion
18  -- has been reached
19  public done : bool := false;
20
21  -- Counter keeping track of the
22  -- simulation progress
23  public runCount : nat := 0;
24
25  -- Number of control periods that have to
26  -- be generated
27  public exLimit : nat := 30;
28
29 operations
30
31 -- Associate a GPIO block to the PWM unit
32 -- controller
33 public setGPIO : gpioInterface ==> ()
34 setGPIO(g) ==
35   outputInt := g;
36
37 -- Associate a controller to the PWMgenerator
38 public PWMunit : Controller ==> PWMunit
39 PWMunit(c) ==
40   controller := c;
41
42 -- Generate a single control cycle. Two different
43 -- control puses can be selected depending on the
44 -- target position.
45 private generateSignal : () ==> ()
46 generateSignal() ==
47 (
48   while runCount < exLimit do
49   (
50     if (getPeriod() = 2) then
```

PWMgenerator

```
51     (-- Set output to high
52     duration (0) outputInt.toggleBit(1);
53     duration (2E6) (); -- Wait 2 milliseconds
54     -- Set output to low
55     duration (0) outputInt.toggleBit(1);
56     duration (18E6) (); -- Wait 18 milliseconds
57     duration (0) runCount := runCount +1;
58     )
59     -- In the second control case 1 millisecond
60     -- is spent in the control pulse. The rest
61     -- of the operation is analogous to the
62     -- previous case.
63     else if getPeriod() = 1 then
64     (
65     duration (0) outputInt.toggleBit(1);
66     duration (1E6) ();
67     duration (0) outputInt.toggleBit(1);
68     duration (19E6) ();
69     duration (0) runCount := runCount +1;
70     )
71     );
72     duration (0) done := true -- Mark the thread as completed
73 );
74
75 -- Get completion flag
76 public getFlag : () ==> bool
77 getFlag() ==
78     return done;
79
80 -- Get the selected period by the discrete processing
81 -- unit
82 private getPeriod : () ==> nat
83 getPeriod() ==
84     return periodSelector;
85
86 -- Sets the control period signal
87 public setPeriod: nat ==> ()
88 setPeriod(p) ==
89     periodSelector := p;
90
91 thread
92     -- Procedural thread, running until completion
93     generateSignal();
94
95 sync
96     -- The signal generation should not be executed
97     -- while the completion flag is requested
98     mutex(generateSignal, getFlag);
99
```

Appendix B. Servo case study VDM models

```
100  -- There should not be two instances of generateSignal
101  -- running in parallel
102  mutex(generateSignal);
103
104  -- Target period must not be set while the period is
105  -- retrieved
106  mutex(setPeriod,getPeriod);
107
108 end PWMunit
```

B.12.5 Sampler

```
1 class Sampler
2
3 instance variables
4
5  -- Sequence of read values
6  public log : seq of nat := [];
7
8  -- Auxiliary variable in order to store the result of
9  -- the write to file operation.
10 writeResult : bool := false;
11
12  -- Instance of the IO class in order to write
13  -- to the file
14  io :IO := new IO();
15
16  -- Skip the first three spurious values
17  public r : nat := 3;
18
19 operations
20
21  -- Toggle the read value in the sampling input.
22  async public toggle : () ==> ()
23  toggle() ==
24  (
25    duration(0) log := log ^ [time];
26  );
27
28  public dumpSignal : () ==> ()
29  dumpSignal () ==
30  (
31    -- Create signal.csv in the case it does not exist
32    -- Clean up previous contents in the case the file exists
33    writeResult := io.fecho("signal.csv", "", <start>);
34
```

gpioInterface

```
35 -- Write the logged time values with the proper format
36 while r < (len Deployment `pwm.outLog.log - 1)
37 do
38 (
39   writeResult :=
40     io.fwriteval[int] ("signal.csv",
41                       Deployment `pwm.outLog.log (r),
42                       <append>);
43   writeResult := io.fecho("signal.csv", ",0", <append>);
44   writeResult := io.fecho("signal.csv", "\n", <append>);
45
46   writeResult :=
47     io.fwriteval[int] ("signal.csv",
48                       Deployment `pwm.outLog.log (r),
49                       <append>);
50   writeResult := io.fecho("signal.csv", ",1", <append>);
51   writeResult := io.fecho("signal.csv", "\n", <append>);
52
53   writeResult :=
54     io.fwriteval[int] ("signal.csv",
55                       Deployment `pwm.outLog.log (r+1),
56                       <append>);
57   writeResult := io.fecho("signal.csv", ",1", <append>);
58   writeResult := io.fecho("signal.csv", "\n", <append>);
59
60   writeResult :=
61     io.fwriteval[int] ("signal.csv",
62                       Deployment `pwm.outLog.log (r+1),
63                       <append>);
64   writeResult := io.fecho("signal.csv", ",0", <append>);
65   writeResult := io.fecho("signal.csv", "\n", <append>);
66
67   r := r+2;
68 )
69 );
70
71 sync
72 -- Two toggle operations cannot be run at the same time
73   mutex(toggle);
74
75 end Sampler
```

B.12.6 gpioInterface

```
1 class gpioInterface
2
```

Appendix B. Servo case study VDM models

```
3 instance variables
4
5   -- Bus width
6   width : nat := 0;
7   -- State of each bit (true = high state, false = low state)
8   state : seq of bool := [];
9
10  -- For each bit there is a state entry in the sequence
11  inv len state = width;
12
13 operations
14
15 -- Class constructor. All states initialized to false by
16 -- default
17 public gpioInterface: nat ==> gpioInterface
18   gpioInterface(w) ==
19   (
20     width := w;
21
22     for all r in set {1,...,width}
23     do state := state ^ [false];
24
25     IO`print(state);
26   )
27 post width = len state and
28     forall i in set {1,...,len state} & state(i) = false;
29
30 -- Set a concrete bit to high level
31 public setHigh: nat ==> ()
32 setHigh(i) ==
33   state(i) := true
34 pre i <= len state; -- Bit position should be in bit sequence
35
36 -- Set a concrete bit to low level
37 public setLow: nat ==> ()
38 setLow(i) ==
39   state(i) := false
40 pre i <= len state; -- Bit position should be in bit sequence
41
42 -- Toggles the state of a certain bit.
43 -- If the bit position is at high level goes to false
44 -- If the bit position is at low level goes to true
45 async public toggleBit: nat ==> ()
46 toggleBit(i) ==
47   (
48     if state(i)
49     then state(i) := not state(i)
50     else state(i) := true;
51     Environment`sampler1.toggle();
```

gpioInterface

```
52 )
53 pre i<=len state; -- Bit position should be in bit sequence
54
55 -- Get the state of a certain bit
56 public getState: nat ==> bool
57 getState(i) ==
58     return state(i)
59 pre i <= len state;
60
61 -- Shows the state of all the bits in the register
62 public showState: () ==> ()
63 showState() ==
64     for all r in set {1,...,width}
65     do showSingle(r);
66
67 -- Get the width of the gpio register
68 public getWidth: () ==> nat
69 getWidth() ==
70     return width;
71
72 -- Prints out the state of a certain bit
73 public showSingle: nat ==> ()
74 showSingle(i) ==
75 (
76     IO`print("\n bit ");
77     IO`print(i);
78     IO`print(": ");
79     IO`print(getState(i));
80 )
81 pre i<=len state; -- Bit position should be in bit sequence
82
83 sync
84     -- Only one operation can modify the state of a bit
85     -- at a certain point of time.
86     mutex(setLow,setHigh);
87
88     -- Bit state cannot be retrieved if it is being set high
89     -- and vice versa.
90     mutex(setHigh,getState);
91
92     -- Bit state cannot be retrieved if it is being set low
93     -- and vice versa.
94     mutex(setLow,getState);
95
96 end gpioInterface
```

B.13. Real-time VDM model: RTOS based solution

B.13.1 CPUloader

```

1 class Controller
2
3 -- Definition of the CPU loader running in a controller
4
5 values
6   -- Execution limit for the CPU loader
7   executionLimit : nat = 90;
8
9 instance variables
10
11   -- Switch time mark. Initialize to a higher value
12   -- in the case a discrete control is used from the
13   -- loader.
14   switchTimeMark :nat := 0;
15
16   progressCounter : nat := 0;
17   pwmref: PWMunit;
18   cSent :bool := false;
19
20 operations
21
22 -- Specify the switch time mark in the case several
23 -- control comands are issued to the controller
24 public setSwitchTimeMark : nat ==> ()
25 setSwitchTimeMark(tm) ==
26   switchTimeMark := tm;
27
28 -- Time consumer operation loading the CPU. The
29 -- function is not performing any computation, just
30 -- taking CPU time.
31 public timeConsumer : () ==> ()
32 timeConsumer () ==
33 (
34   -- Simulated operations performed during
35   -- the first 9 iterations.
36   cases progressCounter:
37     (1) -> duration(1E6) timeEater(),
38     (2) -> duration(2E6) timeEater(),
39     (3) -> duration(3E6) timeEater(),
40     (4) -> duration(4E6) timeEater(),
41     (5) -> duration(5E6) timeEater(),
42     (6) -> duration(6E6) timeEater(),
43     (7) -> duration(7E6) timeEater(),
44     (8) -> duration(8E6) timeEater(),

```


Deployment

```
45     (9) -> duration(9E6) timeEater(),
46     others -> ()
47     end;
48
49     -- Increment the progress counter
50     progressCounter := progressCounter + 1;
51 );
52
53 -- Issues a change in the continuous time hardware
54 -- signal generation.
55 public controlSignal : () ==> ()
56 controlSignal () ==
57 (
58     pwmref.setPeriod(2);
59 );
60
61 -- Establish the reference to the PWM unit
62 public setPWM : PWMunit ==> ()
63 setPWM(p) ==
64     pwmref := p;
65
66 -- Do nothing
67 public timeEater : () ==> ()
68 timeEater() ==
69     skip;
70
71 thread
72     -- Periodic thread executed every 55 milliseconds
73     periodic (55E6,0,0,0) (timeConsumer);
74
75 sync
76     mutex(timeConsumer);
77
78 end Controller
```

B.13.2 Deployment

```
1 system Deployment
2
3 instance variables
4
5     -- Definition of a controller unit, acting as CPU
6     controller : CPU := new CPU(<FP>, 1E9);
7
8     -- Definition of a hardware GPIO block
9     gpioBlock : CPU := new CPU(<FCFS>, 100E6);
```

Appendix B. Servo case study VDM models

```
10
11  -- Associate the controller to the GPIO block
12  controlRegister : BUS := new BUS(<CSMACD>,
13                                72E13,
14                                {controller,gpioBlock});
15
16  -- Static definition of the deployable objects
17  public static loader : Controller := new Controller();
18  public static pwm : PWMunit := new PWMunit();
19  public static gpio : gpioInterface := new gpioInterface(2);
20
21
22
23  operations
24
25  -- Setup and actual deployment
26  public Deployment : () ==> Deployment
27  Deployment () ==
28  (
29    -- Set a reference to the pwm generator from the loader
30    loader.setPWM(pwm);
31    -- Set a reference to the GPIO block from the PWMgenerator
32    pwm.setGPIO(gpio);
33
34
35
36    -- Deploy the active objects in different CPUs
37    controller.deploy(loader);
38    controller.deploy(pwm);
39
40
41    --controller.setPriority(loader.timeConsumer,5);
42
43    gpioBlock.deploy(gpio);
44    controller.setPriority(PWMunit`generateSignal,10);
45    controller.setPriority(Controller`timeConsumer,1);
46
47  );
48
49  end Deployment
```

B.13.3 Environment

```
1  class Environment
2
3  values
```

Environment

```
4
5  -- Time limit mark in the case two control signals
6  -- are going to be generated.
7  private timeLimitMark : nat = 0;
8
9  instance variables
10
11  -- Create a sampler
12  public static sampler1 : Sampler := new Sampler();
13
14  operations
15
16  -- Run operation, simulation entry point
17  public run : () ==> ()
18  run() ==
19  (
20
21  -- Associate the limit time mark for the loader
22  Deployment `loader.setSwitchTimeMark(timeLimitMark);
23
24  -- Associate the sampler to the PWM generator
25  Deployment `pwm.outLog := sampler1;
26
27  -- Start the threads
28  start (Deployment `loader);
29  start (Deployment `pwm);
30
31
32  -- wait(); -- Block until threads are done
33  if not Deployment `pwm.getFlag() then wait();
34
35  -- Dump the sampled signals to a file
36  sampler1.dumpSignal();
37
38  -- Simulation tear down
39  printAndBye();
40
41 );
42
43 -- Simulation tear down
44 public printAndBye: () ==> ()
45 printAndBye() ==
46 (
47   IO `print("#] Generated signal: ");
48   IO `print(Deployment `pwm.outLog.log);
49   IO `print(" [#]");
50   IO `print("\n[1] Model over");
51   IO `print("\n[2] Logs generated");
52 );
```

Appendix B. Servo case study VDM models

```
53 |
54 | -- Wait until it is down
55 | public wait : () ==> ()
56 | wait() ==
57 |     skip;
58 |
59 | sync
60 |     -- Sync predicate for the wait operation
61 |     per wait => Deployment `pwm.getFlag();
62 |
63 | end Environment
```

B.13.4 PWMgenerator

```
1 | class PWMunit
2 |
3 | instance variables
4 |
5 |     -- References to the sampler and the
6 |     -- controller
7 |     public outLog : Sampler;
8 |     public controller : Controller;
9 |
10 |     -- Allows the generation of two different
11 |     -- predefined signals
12 |     private periodSelector : nat := 1;
13 |
14 |     -- Sets a reference to the GPIO interface
15 |     public outputInt : gpioInterface;
16 |
17 |     -- Done flag updated when operation completion
18 |     -- has been reached
19 |     public done : bool := false;
20 |
21 |     -- Counter keeping track of the
22 |     -- simulation progress
23 |     public runCount : nat := 0;
24 |
25 |     -- Number of control periods that have to
26 |     -- be generated
27 |     static public exLimit : nat := 30;
28 |
29 | operations
30 |
31 |     -- Associate a GPIO block to the PWM unit
32 |     -- controller
```

PWMgenerator

```
33 public setGPIO : gpioInterface ==> ()
34 setGPIO(g) ==
35     outputInt := g;
36
37 -- Associate a controller to the PWMgenerator
38 public PWMunit : Controller ==> PWMunit
39 PWMunit(c) ==
40     controller := c;
41
42 public generateSignal : () ==> ()
43 generateSignal () ==
44 (
45     if runCount < exLimit then
46     (
47         duration(2E6) outputInt.toggleBit(1);
48         duration(0) outputInt.toggleBit(1);
49         duration(18E6) ();
50         runCount := runCount + 1;
51     )
52     else
53     (
54         duration (0) done := true;
55     );
56 );
57
58 -- Get completion flag
59 public getFlag : () ==> bool
60 getFlag() ==
61     return done;
62
63 -- Get the selected period by the discrete processing
64 -- unit
65 private getPeriod : () ==> nat
66 getPeriod() ==
67     return periodSelector;
68
69 -- Sets the control period signal
70 public setPeriod: nat ==> ()
71 setPeriod(p) ==
72     periodSelector := p;
73
74 thread
75     -- Periodic invocation to the generate signal generation
76     -- Operation invoked once every 20 milliseconds
77     periodic (20E6,0,0,0) (generateSignal)
78
79
80 sync
81
```

Appendix B. Servo case study VDM models

```
82  mutex(generateSignal, getFlag);
83  mutex(generateSignal);
84  mutex(setPeriod, getPeriod);
85
86 end PWMunit
```

B.13.5 Sampler

```
1  class Sampler
2
3  instance variables
4
5    -- Sequence of read values
6    public log : seq of nat := [];
7
8    -- Auxiliary variable in order to store the result of
9    -- the write to file operation.
10   writeResult : bool := false;
11
12   -- Instance of the IO class in order to write
13   -- to the file
14   io :IO := new IO();
15
16   -- Skip the first three spurious values
17   public r : nat := 3;
18
19 operations
20
21 -- Toggle the read value in the sampling input.
22 async public toggle : () ==> ()
23 toggle() ==
24 (
25   duration(0) log := log ^ [time];
26 );
27
28 public dumpSignal : () ==> ()
29 dumpSignal () ==
30 (
31   -- Create signal.csv in the case it does not exist
32   -- Clean up previous contents in the case the file exists
33   writeResult := io.fecho("signal.csv", "", <start>);
34
35   -- Write the logged time values with the proper format
36   while r < (len Deployment`pwm.outLog.log - 1)
37   do
38   (
```

gpioInterface

```
39   writeResult :=
40     io.fwriteval[int] ("signal.csv",
41       Deployment `pwm.outLog.log(r),
42       <append>);
43   writeResult := io.fecho("signal.csv", ",0", <append>);
44   writeResult := io.fecho("signal.csv", "\n", <append>);
45
46   writeResult :=
47     io.fwriteval[int] ("signal.csv",
48       Deployment `pwm.outLog.log(r),
49       <append>);
50   writeResult := io.fecho("signal.csv", ",1", <append>);
51   writeResult := io.fecho("signal.csv", "\n", <append>);
52
53   writeResult :=
54     io.fwriteval[int] ("signal.csv",
55       Deployment `pwm.outLog.log(r+1),
56       <append>);
57   writeResult := io.fecho("signal.csv", ",1", <append>);
58   writeResult := io.fecho("signal.csv", "\n", <append>);
59
60   writeResult :=
61     io.fwriteval[int] ("signal.csv",
62       Deployment `pwm.outLog.log(r+1),
63       <append>);
64   writeResult := io.fecho("signal.csv", ",0", <append>);
65   writeResult := io.fecho("signal.csv", "\n", <append>);
66
67   r := r+2;
68 )
69 );
70
71 sync
72   -- Two toggle operations cannot be run at the same time
73   mutex(toggle);
74
75 end Sampler
```

B.13.6 gpioInterface

```
1 class gpioInterface
2
3 instance variables
4
5   -- Bus width
6   width : nat := 0;
```

Appendix B. Servo case study VDM models

```

7  -- State of each bit (true = high state, false = low state)
8  state : seq of bool := [];
9
10 -- For each bit there is a state entry in the sequence
11 inv len state = width;
12
13 operations
14
15 -- Class constructor. All states initialized to false by
16 -- default
17 public gpioInterface: nat ==> gpioInterface
18   gpioInterface(w) ==
19   (
20     width := w;
21
22     for all r in set {1,...,width}
23       do state := state ^ [false];
24   IO`print(state);
25   )
26 post width = len state and
27   forall i in set {1,...,len state} & state(i) = false;
28
29 -- Set a concrete bit to high level
30 public setHigh: nat ==> ()
31 setHigh(i) ==
32   state(i) := true
33 pre i <= len state; -- Bit position should be in bit sequence
34
35 -- Set a concrete bit to low level
36 public setLow: nat ==> ()
37 setLow(i) ==
38   state(i) := false
39 pre i <= len state; -- Bit position should be in bit sequence
40
41 -- Toggles the state of a certain bit.
42 -- If the bit position is at high level goes to false
43 -- If the bit position is at low level goes to true
44 async public toggleBit: nat ==> ()
45 toggleBit(i) ==
46   (
47     if state(i)
48     then state(i) := not state(i)
49     else state(i) := true;
50     Environment`sampler1.toggle();
51   )
52 pre i<=len state; -- Bit position should be in bit sequence
53
54 -- Get the state of a certain bit
55 public getState: nat ==> bool

```


gpioInterface

```
56 getState(i) ==
57     return state(i)
58 pre i <= len state;
59
60 -- Shows the state of all the bits in the register
61 public showState: () ==> ()
62 showState() ==
63     for all r in set {1,...,width}
64     do showSingle(r);
65
66 -- Get the width of the gpio register
67 public getWidth: () ==> nat
68 getWidth() ==
69     return width;
70
71 -- Prints out the state of a certain bit
72 public showSingle: nat ==> ()
73 showSingle(i) ==
74 (
75     IO`print("\n bit ");
76     IO`print(i);
77     IO`print(": ");
78     IO`print(getState(i));
79 )
80 pre i<= len state; -- Bit position should be in bit sequence
81
82 sync
83     -- Only one operation can modify the state of a bit
84     -- at a certain point of time.
85     mutex(setLow,setHigh);
86
87     -- Bit state cannot be retrieved if it is being set high
88     -- and vice versa.
89     mutex(setHigh,getState);
90
91     -- Bit state cannot be retrieved if it is being set low
92     -- and vice versa.
93     mutex(setLow,getState);
94
95
96 end gpioInterface
```

B.14. Real-time VDM model: Hardware based solution**B.15. CPUloader**

```

1
2 class Controller
3
4 -- Definition of the CPU loader running in a controller
5
6 values
7   -- Execution limit for the CPU loader
8   executionLimit : nat = 90;
9
10 instance variables
11
12   -- Switch time mark. Initialize to a higher value
13   -- in the case a discrete control is used from the
14   -- loader.
15   switchTimeMark :nat := 0;
16
17   progressCounter : nat := 0;
18   pwmref: PWMunit;
19   cSent :bool := false;
20
21 operations
22
23 -- Specify the switch time mark in the case several
24 -- control comands are issued to the controller
25 public setSwitchTimeMark : nat ==> ()
26 setSwitchTimeMark(tm) ==
27   switchTimeMark := tm;
28
29 -- Time consumer operation loading the CPU. The
30 -- function is not performing any computation, just
31 -- taking CPU time.
32 public timeConsumer : () ==> ()
33 timeConsumer () ==
34 (
35   -- Simulated operations performed during
36   -- the first 9 iterations.
37   cases progressCounter:
38     (1) -> duration(1E6) timeEater(),
39     (2) -> duration(5E6) timeEater(),
40     (3) -> duration(10E6) timeEater(),
41     (4) -> duration(15E6) timeEater(),
42     (5) -> duration(20E6) timeEater(),
43     (6) -> duration(25E6) timeEater(),

```

CPUloader

```
44     (7) -> duration(30E6) timeEater(),
45     (8) -> duration(40E6) timeEater(),
46     (9) -> duration(50E6) timeEater(),
47     others -> ()
48 end;
49
50 -- Increment the progress counter
51 progressCounter := progressCounter + 1;
52
53 -- Send a control signal to the hardware block
54 -- in the case a switchTimeMark has been specified
55 if time > switchTimeMark and not cSent then
56 (
57     controlSignal();
58     IO`print("[#] PWM generator notified at time: ");
59     IO`print(time); IO`print(" [#]\n");
60     cSent := true;
61 );
62 );
63
64 -- Issues a change in the continuous time hardware
65 -- signal generation.
66 public controlSignal : () ==> ()
67 controlSignal () ==
68 (
69     pwmref.setPeriod(2);
70 );
71
72 -- Establish the reference to the PWM unit
73 public setPWM : PWMunit ==> ()
74 setPWM(p) ==
75     pwmref := p;
76
77 -- Do nothing
78 public timeEater : () ==> ()
79 timeEater() ==
80     skip;
81
82 thread
83     -- Procedural thread running unit completion
84     while progressCounter < executionLimit do timeConsumer();
85
86 sync
87     -- Ensures that only one timeConsumer instance is running
88     -- at a certain point of time.
89     mutex(timeConsumer);
90
91 end Controller
```

B.16. Deployment

```

1 system Deployment
2
3 instance variables
4
5   -- Definition of a hardware block for the generation of the PWM
6   -- signal
7   PWMgenerator : CPU := new CPU(<FCFS>, 1E9);
8
9   -- Definition of a controller unit, acting as CPU
10  controller : CPU := new CPU(<FCFS>, 1E5);
11
12  -- Definition of a hardware GPIO block
13  gpioBlock : CPU := new CPU(<FCFS>, 100E6);
14
15  -- Associate the PWMgenerator to the controller
16  -- via control register
17  controlRegister : BUS :=
18    new BUS(<CSMACD>,
19           72E13,
20           PWMgenerator, controller));
21
22  -- Associate the GPIOblock to the PWMgenerator
23  -- via control register
24  controlRegister2 : BUS :=
25    new BUS(<CSMACD>,
26           72E13,
27           {PWMgenerator, gpioBlock});
28
29  -- Static definition of the deployable objects
30  public static loader : Controller := new Controller();
31  public static pwm : PWMunit := new PWMunit();
32  public static gpio : gpioInterface := new gpioInterface(2);
33
34 operations
35
36 -- Setup and actual deployment
37 public Deployment : () ==> Deployment
38 Deployment () ==
39 (
40   -- Set a reference to the pwm generator from the loader
41   loader.setPWM(pwm);
42   -- Set a reference to the GPIO block from the PWMgenerator
43   pwm.setGPIO(gpio);
44
45   -- Deploy the active objects in different CPUs
46   PWMgenerator.deploy(pwm);

```

Environment

```
47 controller.deploy(loader);
48 gpioBlock.deploy(gpio);
49 );
50
51 end Deployment
```

B.17. Environment

```
1 class Environment
2
3 values
4
5     -- Time limit mark in the case two control signals
6     -- are going to be generated.
7     private timeLimitMark : nat = 0;
8
9 instance variables
10
11     -- Create a sampler
12     public static sampler1 : Sampler := new Sampler();
13
14 operations
15
16     -- Run operation, simulation entry point
17     public run : () ==> ()
18     run() ==
19     (
20         -- Associate the limit time mark for the loader
21         Deployment`loader.setSwitchTimeMark(timeLimitMark);
22
23         -- Associate the sampler to the PWM generator
24         Deployment`pwm.outLog := sampler1;
25
26         -- Start the threads
27         start (Deployment`pwm);
28         start (Deployment`loader);
29
30         wait(); -- Block until threads are done
31
32         -- Dump the sampled signals to a file
33         sampler1.dumpSignal();
34
35         -- Simulation tear down
36         printAndBye();
```

```

37
38 );
39
40 -- Simulation tear down
41 public printAndBye: () ==> ()
42 printAndBye() ==
43 (
44   IO`print("[#] Generated signal: ");
45   IO`print(Deployment `pwm.outLog.log);
46   IO`print(" [#]");
47   IO`print("\n[1] Model over");
48   IO`print("\n[2] Logs generated");
49 );
50
51 -- Wait until it is down
52 public wait : () ==> ()
53 wait() ==
54   skip;
55
56 sync
57   -- Sync predicate for the wait operation
58   per wait => Deployment `pwm.getFlag();
59
60 end Environment

```

B.18. PWMgenerator

```

1 class PWMunit
2
3 instance variables
4
5   -- References to the sampler and the
6   -- controller
7   public outLog : Sampler;
8   public controller : Controller;
9
10  -- Allows the generation of two different
11  -- predefined signals
12  private periodSelector : nat := 1;
13
14  -- Sets a reference to the GPIO interface
15  public outputInt : gpioInterface;
16
17  -- Done flag updated when operation completion

```

PWMgenerator

```
18  -- has been reached
19  public done : bool := false;
20
21  -- Counter keeping track of the
22  -- simulation progress
23  public runCount : nat := 0;
24
25  -- Number of control periods that have to
26  -- be generated
27  public exLimit : nat := 30;
28
29  operations
30
31  -- Associate a GPIO block to the PWM unit
32  -- controller
33  public setGPIO : gpioInterface ==> ()
34  setGPIO(g) ==
35    outputInt := g;
36
37  -- Associate a controller to the PWMgenerator
38  public PWMunit : Controller ==> PWMunit
39  PWMunit(c) ==
40    controller := c;
41
42  -- Generate a single control cycle. Two different
43  -- control puses can be selected depending on the
44  -- target position.
45  private generateSignal : () ==> ()
46  generateSignal() ==
47  (
48    while runCount < exLimit do
49    (
50      if (getPeriod() = 2) then
51        ( -- Set output to high
52          duration (0) outputInt.toggleBit(1);
53          duration (2E6) (); -- Wait 2 milliseconds
54          -- Set output to low
55          duration (0) outputInt.toggleBit(1);
56          duration (18E6) (); -- Wait 18 milliseconds
57          duration (0) runCount := runCount +1;
58        )
59      -- In the second control case 1 millisecond
60      -- is spent in the control pulse. The rest
61      -- of the operation is analogous to the
62      -- previous case.
63      else if getPeriod() = 1 then
64        (
65          duration (0) outputInt.toggleBit(1);
66          duration (1E6) ();
```

Appendix B. Servo case study VDM models

```
67     duration (0) outputInt.toggleBit(1);
68     duration (19E6) ();
69     duration (0) runCount := runCount +1;
70   )
71 );
72 duration (0) done := true -- Mark the thread as completed
73 );
74
75 -- Get completion flag
76 public getFlag : () ==> bool
77 getFlag() ==
78   return done;
79
80 -- Get the selected period by the discrete processing
81 -- unit
82 private getPeriod : () ==> nat
83 getPeriod() ==
84   return periodSelector;
85
86 -- Sets the control period signal
87 public setPeriod: nat ==> ()
88 setPeriod(p) ==
89   periodSelector := p;
90
91 thread
92   -- Procedural thread, running until completion
93   generateSignal();
94
95 sync
96   -- The signal generation should not be executed
97   -- while the completion flag is requested
98   mutex(generateSignal,getFlag);
99
100   -- There should not be two instances of generateSignal
101   -- running in parallel
102   mutex(generateSignal);
103
104   -- Target period must not be set while the period is
105   -- retrieved
106   mutex(setPeriod,getPeriod);
107
108 end PWMunit
```

B.19. Sampler

Sampler

```
1 class Sampler
2
3 instance variables
4
5   -- Sequence of read values
6   public log : seq of nat := [];
7
8   -- Auxiliary variable in order to store the result of
9   -- the write to file operation.
10  writeResult : bool := false;
11
12  -- Instance of the IO class in order to write
13  -- to the file
14  io :IO := new IO();
15
16  -- Skip the first three spurious values
17  public r : nat := 3;
18
19 operations
20
21 -- Toggle the read value in the sampling input.
22 async public toggle : () ==> ()
23 toggle() ==
24 (
25   duration(0) log := log ^ [time];
26 );
27
28 public dumpSignal : () ==> ()
29 dumpSignal () ==
30 (
31   -- Create signal.csv in the case it does not exist
32   -- Clean up previous contents in the case the file exists
33   writeResult := io.fecho("signal.csv", "", <start>);
34
35   -- Write the logged time values with the proper format
36   while r < (len Deployment `pwm.outLog.log - 1)
37   do
38   (
39     writeResult :=
40       io.fwriteval[int]("signal.csv",
41         Deployment `pwm.outLog.log(r),
42         <append>);
43     writeResult := io.fecho("signal.csv", ",0", <append>);
44     writeResult := io.fecho("signal.csv", "\n", <append>);
45
46     writeResult :=
47       io.fwriteval[int]("signal.csv",
48         Deployment `pwm.outLog.log(r),
```

Appendix B. Servo case study VDM models

```
49         <append>);
50     writeResult := io.fecho("signal.csv", ",1", <append>);
51     writeResult := io.fecho("signal.csv","\n", <append>);
52
53     writeResult :=
54         io.fwriteval[int]("signal.csv",
55             Deployment `pwm.outLog.log(r+1),
56             append>);
57     writeResult := io.fecho("signal.csv", ",1", <append>);
58     writeResult := io.fecho("signal.csv","\n", <append>);
59
60     writeResult :=
61         io.fwriteval[int]("signal.csv",
62             Deployment `pwm.outLog.log(r+1),
63             <append>);
64     writeResult := io.fecho("signal.csv", ",0", <append>);
65     writeResult := io.fecho("signal.csv","\n", <append>);
66
67     r := r+2;
68 )
69 );
70
71 sync
72     -- Two toggle operations cannot be run at the same time
73     mutex(toggle);
74
75 end Sampler
```

B.20. gpioInterface

```
1 class gpioInterface
2
3 instance variables
4
5     -- Bus width
6     width : nat := 0;
7     -- State of each bit (true = high state, false = low state)
8     state : seq of bool := [];
9
10    -- For each bit there is a state entry in the sequence
11    inv len state = width;
12
13 operations
14
```

gpioInterface

```

15  -- Class constructor. All states initialized to false by
16  -- default
17  public gpioInterface: nat ==> gpioInterface
18    gpioInterface(w) ==
19    (
20      width := w;
21
22      for all r in set {1,...,width}
23        do state := state ^ [false];
24
25      IO`print(state);
26
27    )
28  post width = len state and
29    forall i in set {1,...,len state} & state(i) = false;
30
31  -- Set a concrete bit to high level
32  public setHigh: nat ==> ()
33  setHigh(i) ==
34    state(i) := true
35  pre i <= len state; -- Bit position should be in bit sequence
36
37  -- Set a concrete bit to low level
38  public setLow: nat ==> ()
39  setLow(i) ==
40    state(i) := false
41  pre i <= len state; -- Bit position should be in bit sequence
42
43  -- Toggles the state of a certain bit.
44  -- If the bit position is at high level goes to false
45  -- If the bit position is at low level goes to true
46  async public toggleBit: nat ==> ()
47  toggleBit(i) ==
48  (
49    if state(i)
50    then state(i) := not state(i)
51    else state(i) := true;
52    Environment`sampler1.toggle();
53  )
54  pre i<=len state; -- Bit position should be in bit sequence
55
56  -- Get the state of a certain bit
57  public getState: nat ==> bool
58  getState(i) ==
59    return state(i)
60  pre i <= len state;
61
62  -- Shows the state of all the bits in the register
63  public showState: () ==> ()

```

Appendix B. Servo case study VDM models

```
64 showState() ==
65   for all r in set {1,...,width}
66   do showSingle(r);
67
68 -- Get the width of the gpio register
69 public getWidth: () ==> nat
70 getWidth() ==
71   return width;
72
73 -- Prints out the state of a certain bit
74 public showSingle: nat ==> ()
75 showSingle(i) ==
76 (
77   IO`print("\n bit ");
78   IO`print(i);
79   IO`print(": ");
80   IO`print(getState(i));
81 )
82 pre i<= len state; -- Bit position should be in bit sequence
83
84 sync
85   -- Only one operation can modify the state of a bit
86   -- at a certain point of time.
87   mutex(setLow,setHigh);
88
89   -- Bit state cannot be retrieved if it is being set high
90   -- and vice versa.
91   mutex(setHigh,getState);
92
93   -- Bit state cannot be retrieved if it is being set low
94   -- and vice versa.
95   mutex(setLow,getState);
96
97 end gpioInterface
```

AVB case study VDM models

This appendix contains the VDM models for the AVB case study. Section C.1 presents the initial overall sequential VDM models. Section C.2 presents the sequential VDM models. Section C.3 presents the concurrent VDM models. Finally, section C.4 presents the real-time VDM models.

C.1. Overall sequential model

C.1.1 App

```

1 class App
2
3 instance variables
4
5     private preciseOriginTS : real;
6     private correctionField : real;
7
8     private localTimeStamp : real;
9
10
11 operations
12
13 public receiveTimeInformation : seq of real ==> ()
14 receiveTimeInformation (time) ==
15 (
16     preciseOriginTS := time(1);
17     correctionField := time(2);
18     localTimeStamp := time(3);
19 );
20
21 public showReceivedTime : () ==> ()
22 showReceivedTime() ==
23 (

```

Appendix C. AVB case study VDM models

```
24 IO`print("\nAt time: ");
25 IO`print(localTimeStamp);
26 IO`print("\nThe Grand Master time was: ");
27 IO`print(preciseOriginTS);
28 IO`print("\nIt has been corrected with: ");
29 IO`print(correctionField);
30 IO`print("\nFinal time for previos timestamp is: ");
31 IO`print(preciseOriginTS+correctionField);
32 );
33
34 end App
```

C.1.2 Clock

```
1 class Clock
2
3 instance variables
4
5     time : nat := 0;
6
7 operations
8
9 public tick: () ==> ()
10 tick() ==
11 (
12     time := time +1;
13 );
14
15 public getTime: () ==> nat
16 getTime() ==
17     return time;
18
19 end Clock
```

C.1.3 ClockSlave

```
1 class ClockSlave
2
3 instance variables
4
5     private siteSync : Device;
6
7     private preciseOriginTS : real;
```

ClockSlave

```
8   private correctionField : real;  
9   private localTimeStamp : real;  
10  
11  private timeConsumers : seq of App := [];  
12  
13  operations  
14  
15  public registerApp : App ==> nat  
16  registerApp(ap) ==  
17  (  
18    timeConsumers := timeConsumers ^ [ap];  
19    return len timeConsumers;  
20  );  
21  
22  public setTimeInfo: seq of real ==> ()  
23  setTimeInfo(time) ==  
24  (  
25    preciseOriginTS := time(1);  
26    correctionField := time(2);  
27    localTimeStamp := time(3);  
28    notifyTime();  
29    IO`print("\nSet time information");  
30    IO`print(preciseOriginTS);  
31    IO`print("\n");  
32    IO`print(correctionField);  
33    IO`print("\n");  
34    IO`print(localTimeStamp);  
35  );  
36  
37  public notifyTime : () ==> ()  
38  notifyTime() ==  
39    let t : seq of real =  
40    [preciseOriginTS,correctionField,localTimeStamp] in  
41    for all i in set {1,...,len timeConsumers} do  
42    timeConsumers(i).receiveTimeInformation(t);  
43  
44  private setSiteSync : Device ==> ()  
45  setSiteSync(dv) ==  
46    siteSync := dv;  
47  
48  public getTime:() ==> ()  
49  getTime() ==  
50  ();  
51  
52  public ClockSlave: () ==> ClockSlave  
53  ClockSlave() ==  
54  (  
55    let i : nat = registerApp(new App()) in ();  
56  );
```

```

57
58 public triggerPrint: () ==> ()
59 triggerPrint() ==
60   for all i in set {1,...,len timeConsumers} do
61     timeConsumers(i).showReceivedTime();
62
63 end ClockSlave

```

C.1.4 Device

```

1 class Device
2
3 instance variables
4
5   -- Device ID number
6   private deviceID : nat := 0;
7
8   public isBridge : bool := true;
9
10  -- Frequency of the current device.
11  private freq : nat := 0;
12
13  -- Master ports associated to the device
14  public masterPort : Master := new Master();
15
16  -- Master ports attached to the device
17  public masterPorts : seq of Master := [];
18  inv isBridge <=> len masterPorts >= 1;
19
20  -- Slave port associated to the device
21  public slavePort : Slave := new Slave();
22
23  -- Device with no external slaves connected
24  -- (finalDevice = true)
25  public finalDevice : bool := false;
26  --inv finalDevice => len masterPorts = 0;
27
28  -- Current device is Grand Master -> iGM = true
29  -- if device isGM it should not have masters attached.
30  public isGM : bool := false;
31  inv isGM => not finalDevice;
32  inv finalDevice => not isGM;
33
34  -- Frequency of the master device
35  private masterFreq : nat := 0;
36

```


Device

```
37 private receivedSyncEvent : bool := false;  
38 private syncRecTS : nat := 0;  
39  
40 -- Neighbour rate ratio measured by the slave port  
41 private neighbourRateRatio : real := 1;  
42  
43 -- Precise origin timestamp, defined by GM and  
44 -- never changed.  
45 private preciseOriginTS : real := 0;  
46  
47 -- received correction field, to be recomputed  
48 -- before use.  
49 private receivedCorrectionField : real := 0;  
50  
51 -- CumulativeRateRatio received from the previous  
52 -- node.  
53 private cummulativeRateRatio : real := 1;  
54 -- Delay from master to slave port  
55 private linkDelay : real := 10;  
56 -- Time that has been spent in the current node  
57 private residenceTime : real := 0;  
58  
59 -- Correction field calculated at the current node,  
60 -- to be used in the current node.  
61 private newCorrectionField : real := 0;  
62 -- CummulativeRateRatio calculated at the current node,  
63 -- to be used in the current node.  
64 private newCummulativeRateRatio : real := 0;  
65  
66 -- Clock attached to the device  
67 private clk : Clock;  
68  
69 private timeInfoAvailable : bool;  
70  
71 public cSlave : ClockSlave := new ClockSlave();  
72  
73 operations  
74  
75 public notifyClock : () ==> ()  
76 notifyClock() ==  
77   cSlave.setTimeInfo([preciseOriginTS,  
78                       receivedCorrectionField,  
79                       syncRecTS]);  
80  
81 public calculateResidenceTime : () ==> ()  
82 calculateResidenceTime() ==  
83   if not finalDevice and not isGM then  
84     residenceTime := masterPorts(1).getEgresTime()  
85     - slavePort.getIngressTime()
```

Appendix C. AVB case study VDM models

```

86 pre masterPorts(1).getEgressTime() > slavePort.getIngressTime()
87 post residenceTime > 0;
88
89 public makeTimeAvailable : () ==> ()
90 makeTimeAvailable() ==
91     timeInfoAvailable := true;
92
93 public getTimeInfo : () ==> seq of real
94 getTimeInfo() ==
95     return [preciseOriginTS,
96             receivedCorrectionField,
97             cummulativeRateRatio];
98
99 public timeInfoRead : () ==> ()
100 timeInfoRead() ==
101     timeInfoAvailable := false;
102
103 -- Registers a new master port in the siteSync entity
104 public registerMasterPort : Master ==> ()
105 registerMasterPort(m) ==
106     masterPorts := masterPorts ^ [m];
107
108 -- Sets a single master port for the current device
109 public setMasterPort: Master ==> ()
110 setMasterPort(m) ==
111     masterPort := m;
112
113 -- Sets the slave port that corresponds to the current device
114 public setSlavePort: Slave ==> ()
115 setSlavePort(s) ==
116     slavePort := s;
117
118 -- Sets device ID
119 public setDeviceID: nat ==> ()
120 setDeviceID(id) ==
121     deviceID := id
122 pre deviceID = 0
123 post deviceID <> 0;
124
125 -- Gets device ID
126 public getDeviceID: () ==> nat
127 getDeviceID() ==
128     return deviceID
129 pre deviceID <> 0;
130
131 -- Sets the residence time. Time that will take the system
132 -- to process an incoming sync event.
133 public setResidenceTime : nat ==> ()
134 setResidenceTime(t) ==

```

Device

```
135     residenceTime := t
136 post residenceTime <> 0;
137
138 -- Generates the POTS
139 public generatePOTS : nat ==> ()
140 generatePOTS(ts) ==
141     preciseOriginTS := ts
142 pre isGM
143 post preciseOriginTS <> 0;
144
145 -- Forwards the POTS
146 public sendPOTS : () ==> ()
147 sendPOTS() ==
148     if not finalDevice then
149         masterPort.sendPOTS(preciseOriginTS);
150
151 -- Receives the POTS
152 public receivePOTS: () ==> ()
153 receivePOTS() ==
154     preciseOriginTS := slavePort.receivePOTS();
155
156 -- Sets the frequency of the master device.
157 public setMasterFreq : nat ==> ()
158 setMasterFreq (f) ==
159     masterFreq := f
160 pre f <> 0;
161
162 -- Device constructor
163 public Device : nat ==> Device
164 Device (f) ==
165 (
166     freq := f;
167     registerMasterPort(masterPort);
168 );
169
170 -- Announces the device frequency
171 public announceFreq : () ==> ()
172 announceFreq() ==
173     if (not finalDevice) then
174         for all i in set {1,...,len masterPorts} do
175             masterPorts(i).announceFreq(freq);
176
177 -- Get the frequency of the master device attached to
178 -- the slave port.
179 public getFreq : () ==> ()
180 getFreq() ==
181     if (not isGM) then setMasterFreq(slavePort.getFreq());
182
183 -- Sends the follow up message to all the slave unit
```

Appendix C. AVB case study VDM models

```

184 -- attached to it.
185 -- Transmission process is performed in one operation.
186 public sendFollowUpMessage: () ==> ()
187 sendFollowUpMessage() ==
188   if not finalDevice then
189     for all i in set {1,...,len masterPorts} do
190       masterPorts(i).sendFollowUpMessage(newCorrectionField,
191                                           preciseOriginTS,
192                                           newCummulativeRateRatio);
193
194 -- Sends the cummulative rate ratio to the slave
195 -- unit attached to it.
196 public sendCRR: () ==> ()
197 sendCRR() ==
198   if (not finalDevice) then
199     masterPort.sendCorrectionField(newCummulativeRateRatio);
200
201 -- Receives the follow up message sent by the master.
202 -- Reception is performed in one operation.
203 public receiveFollowUpMessage: () ==> ()
204 receiveFollowUpMessage() ==
205   let s : seq of real = slavePort.receiveFollowUpMessage() in
206   (
207     receivedCorrectionField := s(1);
208     preciseOriginTS := s(2);
209     cummulativeRateRatio := s(3);
210   );
211
212 -- Receives the cummulative rate ratio at the slave
213 -- port. Cummulative rate ratio is sent by the master
214 -- device attached to the slave port.
215 public receiveCRR: () ==> ()
216 receiveCRR() ==
217   cummulativeRateRatio:= slavePort.receiveCRR();
218
219 -- Sets the current device as the grand master
220 -- for the network.
221 public setGM: () ==> ()
222 setGM() ==
223   isGM := true;
224
225 -- Calculates the neighbour rate ratio. This calculation
226 -- is performed at every slave device.
227 public calculateNeighbourRR : () ==> ()
228 calculateNeighbourRR () ==
229   if (not isGM) then
230     neighbourRateRatio := freq / masterFreq
231   else
232     neighbourRateRatio := 1

```

Device

```
233 pre not isGM => (freq <> 0 and masterFreq <> 0)
234 post neighbourRateRatio <> 0;
235
236 -- Calculates the new cummulative rate ratio to be sent
237 -- to the attached slave unit.
238 public calculateCRateRatio : () ==> ()
239 calculateCRateRatio () ==
240     newCummulativeRateRatio := cummulativeRateRatio *
241                               neighbourRateRatio
242 pre neighbourRateRatio <> 0 and cummulativeRateRatio <> 0;
243
244 -- Calculates the new correction field to be sent to the
245 -- attached slave unit.
246 public calculateNewCorrectionField : () ==> ()
247 calculateNewCorrectionField () ==
248     if isGM then ()
249     else (
250         newCorrectionField := linkDelay + residenceTime *
251                               newCummulativeRateRatio + receivedCorrectionField);
252
253 -- Sends a correction field to the slave unit.
254 public sendNewCorrectionField : () ==> ()
255 sendNewCorrectionField () ==
256     if (not finalDevice) then
257         masterPort.sendCorrectionField(newCorrectionField);
258
259 -- Receives a correction field sent by master unit.
260 public receiveCorrectionField : () ==> ()
261 receiveCorrectionField() ==
262     receivedCorrectionField := slavePort.receiveCorrectionField();
263
264 -- Sets the current device as a final device.
265 -- A final device will not have slave devices and
266 -- will not be acting as a master device.
267 public setFinalDevice: () ==> ()
268 setFinalDevice() ==
269     finalDevice := true
270 pre not isGM
271 post finalDevice;
272
273 -- Receives a synchronization sync event and
274 -- time stamps it.
275 public receiveSyncEvent: () ==> ()
276 receiveSyncEvent () ==
277 (
278     receivedSyncEvent := slavePort.receiveSyncEvent();
279     updateTimeRef();
280 );
281
```

Appendix C. AVB case study VDM models

```
282 -- Updates the time reference when the time stamp operation
283 -- is invoked.
284 public updateTimeRef: () ==> ()
285 updateTimeRef() ==
286 (
287   IO`print("\n Sync event --> timestamp generated");
288   -- syncRecTS := clk.getTime();
289 );
290
291 -- Sends a synchronization event to a slave unit.
292 public sendSyncEvent: () ==> ()
293 sendSyncEvent () ==
294   if (not finalDevice) then
295     masterPort.sendSyncEvent();
296
297
298 -- Shows the device state
299 public showDevice : () ==> ()
300 showDevice () ==
301 (
302   IO`print("\nInfo for device: ");
303   IO`print(deviceID);
304   IO`print("\nIs GM:");
305   IO`print(isGM);
306   IO`print("\nIs final device: ");
307   IO`print(finalDevice);
308   IO`print("\nIs sync event received: ");
309   IO`print(receivedSyncEvent);
310   IO`print("\nSync event received at: ");
311   IO`print(syncRecTS);
312   IO`print("\nFrequency: ");
313   IO`print(freq);
314   IO`print("\nMaster frequency: ");
315   IO`print(masterFreq);
316   IO`print("\nLinkDelay: ");
317   IO`print(linkDelay);
318   IO`print("\nResidence Time: ");
319   IO`print(residenceTime);
320   IO`print("\nNeighbour Rate Ratio: ");
321   IO`print(neighbourRateRatio);
322   IO`print("\nRecCummulative Rate Ratio: ");
323   IO`print(cummulativeRateRatio);
324   IO`print("\nNew Cumulative Rate Ratio: ");
325   IO`print(newCummulativeRateRatio);
326   IO`print("\nReceived CorrectionField: ");
327   IO`print(receivedCorrectionField);
328   IO`print("\nNew calculated CorrectionField: ");
329   IO`print(newCorrectionField);
330   IO`print("\nPrecise Origin TimeStamp: ");
```

Main

```
331 IO`print (preciseOriginTS);
332
333 --cSlave.triggerPrint ();
334 );
335
336 end Device
```

C.1.5 Main

```
1 class Main
2
3 instance variables
4
5   clock : Clock := new Clock ();
6
7   maxSimTime : nat := 0;
8
9   -- Devices use the simulation
10  device1 : Device := new Device(30);
11  device2 : Device := new Device(10);
12  device3 : Device := new Device(20);
13  device4 : Device := new Device(10);
14
15  device12 : Device := new Device(20);
16
17  -- Channels used to cummunicate the devices
18  ch1_2 : PhyMessage := new PhyMessage ();
19  ch2_3 : PhyMessage := new PhyMessage ();
20  ch3_4 : PhyMessage := new PhyMessage ();
21
22  ch1_12 : PhyMessage := new PhyMessage ();
23
24  testMP : Master := new Master(ch1_12);
25
26  -- Sequence of devices involved in the scenario
27  deployment : seq of Device := [device1,device2,device3,device4];
28
29  -- Sequence of channels between the devices
30  channels : seq of PhyMessage := [ch1_2,ch2_3,ch3_4];
31
32 operations
33
34 public run : () ==> ()
35 run () ==
36 (
37
```

Appendix C. AVB case study VDM models

```
38  -- Sets up the network for the simulation
39  setupNetwork();
40
41  -- Device 1 operations
42
43  device1.generatePOTS(10);
44  device1.sendSyncEvent();
45
46  device1.calculateCRateRatio();
47  device1.calculateNewCorrectionField();
48  device1.sendFollowUpMessage();
49
50  device12.receiveSyncEvent();
51
52  device12.receiveFollowUpMessage();
53  device12.calculateCRateRatio();
54  device12.calculateNewCorrectionField();
55
56
57  -- Device 2 operations
58
59  device2.receiveSyncEvent();
60  device2.sendSyncEvent();
61
62  device2.receiveFollowUpMessage();
63
64  device2.calculateResidenceTime();
65  device2.calculateCRateRatio();
66  device2.calculateNewCorrectionField();
67
68  device2.sendFollowUpMessage();
69
70  -- Device 3 operations
71
72  device3.receiveSyncEvent();
73  device3.sendSyncEvent();
74
75  device3.receiveFollowUpMessage();
76
77  device3.notifyClock();
78  device3.cSlave.triggerPrint();
79
80  device3.calculateResidenceTime();
81  device3.calculateCRateRatio();
82  device3.calculateNewCorrectionField();
83
84  device3.sendFollowUpMessage();
85
86  -- Device 4 operations
```


Main

```
87
88 device4.receiveSyncEvent();
89 device4.sendSyncEvent();
90
91 device4.receiveFollowUpMessage();
92 device4.calculateCRateRatio();
93 device4.calculateNewCorrectionField();
94
95 device4.sendFollowUpMessage();
96
97 -- End of model
98
99 showDevices();
100 device12.showDevice();
101
102 );
103
104 public setupNetwork: () ==> ()
105 setupNetwork() ==
106 (
107   for all i in set {1,...,len deployment}
108   do
109     deployment(i).setDeviceID(i);
110
111   deployment(1).setGM();
112   deployment(len deployment).setFinalDevice();
113   device12.setFinalDevice();
114
115   for all i in set {1,...,len deployment - 1}
116   do
117     (
118       deployment(i).masterPort.changePhyChannel(channels(i));
119       deployment(i+1).slavePort.changePhyChannel(
120         deployment(i).masterPorts(1).getPhyChannel()
121       );
122     );
123
124   device1.registerMasterPort(testMP);
125   device12.slavePort.changePhyChannel(
126     device1.masterPorts(2).getPhyChannel()
127   );
128   device12.setResidenceTime(10);
129
130   -- Execute frequency announcement and
131   -- neighbour rate ratio calculations
132   -- for the daisy chained devices.
133   announceFreqs();
134   calculateNeighbourRR();
135
```

Appendix C. AVB case study VDM models

```
136  -- Execute frequency announcement and
137  -- neighbour rate ratio calculations
138  -- for the branched devices
139
140  device12.announceFreq();
141  device12.getFreq();
142  device12.calculateNeighbourRR();
143
144 );
145
146 -- Set of actions for the daisy chained devices.
147 public announceFreqs : () ==> ()
148 announceFreqs () ==
149   for all i in set {1,...,len deployment}
150   do
151     (
152       deployment(i).announceFreq();
153       deployment(i).getFreq();
154     );
155
156
157 public calculateNeighbourRR : () ==> ()
158 calculateNeighbourRR () ==
159   for all i in set {1,...,len deployment}
160   do deployment(i).calculateNeighbourRR();
161
162 public showDevices : () ==> ()
163 showDevices () ==
164   for all i in set {1,...,len deployment}
165   do deployment(i).showDevice();
166
167
168 end Main
```

C.1.6 Master

```
1 class Master is subclass of Port
2
3 instance variables
4
5   egressTime : real := 12;
6
7 operations
8
9 public timeStamp : () ==> ()
10 timeStamp () ==
```

Master

```
11 | egressTime := 10;
12 |
13 | public getEgressTime : () ==> real
14 | getEgressTime () ==
15 |     return egressTime;
16 |
17 | -- Class constructor.
18 | -- Receives the physical layer associated to the device
19 | public Master : PhyMessage ==> Master
20 | Master(phy) ==
21 |     channel := phy;
22 |
23 | -- Changes the physical channel associated to the device
24 | public changePhyChannel : PhyMessage ==> ()
25 | changePhyChannel(ch) ==
26 |     channel := ch;
27 |
28 | -- Sends a precise origin time-stamp
29 | public sendPOTS : real ==> ()
30 | sendPOTS(pots) ==
31 |     channel.setPOTS(pots);
32 |
33 | -- Sends a synchronization event
34 | public sendSyncEvent : () ==> ()
35 | sendSyncEvent() ==
36 |     channel.signalSyncEvent();
37 |
38 | -- Sends the cummulative rate ratio
39 | public sendCRR : real ==> ()
40 | sendCRR(crr) ==
41 |     channel.setCRR(crr);
42 |
43 | -- Sends the correction field
44 | public sendCorrectionField : real ==> ()
45 | sendCorrectionField(crf) ==
46 |     channel.setCorrectionField(crf);
47 |
48 | -- Sends the followup up message in one call
49 | -- Follow up messsage is composed by Precise Origing Time Stamp,
50 | -- cummulative rate ratio and correction
51 | -- field in one invocation
52 | public sendFollowUpMessage: real * real * real ==> ()
53 | sendFollowUpMessage(crf,pots, crr) ==
54 |     channel.setFollowUp(crf,pots,crr);
55 |
56 | -- Announce the frequency of the master device controlling
57 | -- the port
58 | public announceFreq: nat ==> ()
59 | announceFreq(f) ==
```

```

60   channel.setMasterFreq(f);
61
62 end Master

```

C.1.7 PhyMessage

```

1 class PhyMessage
2
3 instance variables
4
5   -- sync event flag. Holds true value if a sync event
6   -- is put on the bus.
7   private syncEvent : bool := false;
8
9   -- Contents of the follow up message
10  private preciseOriginTS : real := 0.0;
11  private cumulativeRateRatio : real := 0.0;
12  private correctionField : real := 0.0;
13
14  -- Frequency sent by a master
15  private masterFreq : real := 0;
16
17 operations
18
19  -- Gets the contents of the follow up message.
20  public getFollowUp: () ==> seq of real
21  getFollowUp() ==
22    return [correctionField,
23            preciseOriginTS,
24            cumulativeRateRatio];
25
26  -- Sets the contents of the follow up message.
27  public setFollowUp: real * real * real ==> ()
28  setFollowUp(correction, pots, crr) ==
29  (
30    setCRR(crr);
31    setCorrectionField(correction);
32    setPOTS(pots);
33  );
34
35  -- Signals a Synchronization event
36  public signalSyncEvent: () ==> ()
37  signalSyncEvent() ==
38  (
39    IO`print("\nMaster port --> Sync event sent");
40    syncEvent := true;

```

PhyMessage

```
41 );
42
43 -- Check if there is a Synchronization event
44 public getSyncEvent: () ==> bool
45 getSyncEvent() ==
46 (
47     IO`print("\nSlave port --> Sync event received");
48     return syncEvent;
49 );
50
51 -- #####
52 -- Getters and setters for the different fields in the
53 -- Physical Message.
54
55 public setMasterFreq : real ==> ()
56 setMasterFreq(f) ==
57     masterFreq := f;
58
59 public getMasterFreq: () ==> real
60 getMasterFreq() ==
61     return masterFreq;
62
63 public setCRR: real ==> ()
64 setCRR(crr) ==
65     cummulativeRateRatio := crr;
66
67 public setCorrectionField: real ==> ()
68 setCorrectionField(crf) ==
69     correctionField := crf;
70
71 public setPOTS: real ==> ()
72 setPOTS(pots) ==
73     preciseOriginTS := pots;
74
75 public getCField: () ==> real
76 getCField() ==
77     return correctionField;
78
79 public getPOTS: () ==> real
80 getPOTS() ==
81     return preciseOriginTS;
82
83 public getCRR: () ==> real
84 getCRR() ==
85     return cummulativeRateRatio;
86
87 sync
88
89     mutex(getFollowUp, setFollowUp);
```

```

90   mutex(signalSyncEvent, getSyncEvent);
91
92 end PhyMessage

```

C.1.8 Port

```

1 class Port
2
3 instance variables
4
5   -- Physical channel associated to the port
6   protected channel : PhyMessage;
7
8   -- Associated ID to the port
9   protected portID : nat;
10
11  protected clk : Clock;
12
13 operations
14
15 public timeStamp : () ==> ()
16 timeStamp() ==
17   is not yet specified;
18
19   -- Change the associated channel to the slave
20 public changePhyChannel : PhyMessage ==> ()
21 changePhyChannel(ch) ==
22   channel := ch;
23
24 public getPhyChannel : () ==> PhyMessage
25 getPhyChannel() ==
26   return channel;
27
28   -- Changes the associated ID to the port
29 public changePortID: nat ==> ()
30 changePortID(id) ==
31   portID := id;
32
33 end Port

```

C.1.9 Slave

```

1 class Slave is subclass of Port

```

Slave

```
2
3 instance variables
4
5   ingressTime : real := 2;
6
7 operations
8
9 public timeStamp : () ==> ()
10 timeStamp() ==
11   ingressTime := 2;
12
13 public getIngressTime : () ==> real
14 getIngressTime () ==
15   return ingressTime;
16
17 -- Class constructor
18 public Slave: PhyMessage ==> Slave
19 Slave(phy) ==
20   channel := phy;
21
22 -- Receives the cumulative rate ratio
23 public receiveCRR: () ==> real
24 receiveCRR() ==
25   channel.getCRR();
26
27 -- Receives the correction field
28 public receiveCorrectionField: () ==> real
29 receiveCorrectionField() ==
30   channel.getCField();
31
32 -- Receives the precise origin time stamp
33 public receivePOTS : () ==> real
34 receivePOTS() ==
35   channel.getPOTS();
36
37 -- Receives the follow up message
38 -- Receives Precise Origing Time Stamp, cumulative rate ratio
39 -- and correction field in one invocation
40 public receiveFollowUpMessage: () ==> seq of real
41 receiveFollowUpMessage() ==
42   channel.getFollowUp();
43
44 -- Checks if there is a synchronization event
45 -- available
46 public receiveSyncEvent: () ==> bool
47 receiveSyncEvent() ==
48   channel.getSyncEvent();
49
50 -- Gets the frequency associated to the master
```

```

51 public getFreq: () ==> real
52 getFreq() ==
53     return channel.getMasterFreq();
54
55 end Slave

```

C.2. Sequential model

C.2.1 App

```

1 class App
2
3 instance variables
4
5     private preciseOriginTS : real;
6     private correctionField : real;
7
8     private localTimeStamp : real;
9
10    private timeInfoReady : bool := false;
11
12
13 operations
14
15 public receiveTimeInformation : seq of real ==> ()
16 receiveTimeInformation (time) ==
17 (
18     preciseOriginTS := time(1);
19     correctionField := time(2);
20     localTimeStamp := time(3);
21
22     timeInfoReady := true;
23 );
24
25 public showReceivedTime : () ==> ()
26 showReceivedTime() ==
27 (
28     if timeInfoReady then
29     (
30         IO`print("\nAt time: ");
31         IO`print(localTimeStamp);
32         IO`print("\nThe Grand Master time was: ");
33         IO`print(preciseOriginTS);
34         IO`print("\nIt has been corrected with: ");

```


Clock

```
35 IO`print(correctionField);
36 IO`print("\nFinal time for previos timestamp is: ");
37 IO`print(preciseOriginTS+correctionField);
38 )
39 else
40     IO`print("\n Message from application:
41             Time information not Ready");
42 );
43
44 end App
```

C.2.2 Clock

```
1 class Clock
2
3 instance variables
4
5     time : nat := 0;
6
7 operations
8
9 public tick: () ==> ()
10 tick() ==
11 (
12     time := time +1;
13 );
14
15 public getTime: () ==> nat
16 getTime() ==
17     return time;
18
19 end Clock
```

C.2.3 ClockSlave

```
1 class ClockSlave
2
3 instance variables
4
5     private siteSync : Device;
6
7     private preciseOriginTS : real;
8     private correctionField : real;
```

```

9   private localTimeStamp : real;
10
11  private timeConsumers : seq of App := [];
12
13  private newTimeInfoReady : bool := false;
14
15  operations
16
17  public registerApp : App ==> ()
18  registerApp(ap) ==
19    timeConsumers := timeConsumers ^ [ap];
20
21  public setTimeInfo: seq of real ==> ()
22  setTimeInfo(time) ==
23    (
24    if len time > 2 then
25      (
26        preciseOriginTS := time(1);
27        correctionField := time(2);
28        localTimeStamp := time(3);
29        newTimeInfoReady := true;
30      )
31    );
32
33  public timeStep: () ==> ()
34  timeStep() ==
35    if newTimeInfoReady then
36      (
37        notifyTime();
38        newTimeInfoReady := false
39      )
40    else
41      skip;
42
43  public notifyTime : () ==> ()
44  notifyTime() ==
45    let t : seq of real =
46      [preciseOriginTS,correctionField,localTimeStamp] in
47    for all i in set {1,...,len timeConsumers} do
48      timeConsumers(i).receiveTimeInformation(t);
49
50  public setSiteSync : Device ==> ()
51  setSiteSync(dv) ==
52    siteSync := dv;
53
54  public getTime:() ==> ()
55  getTime() ==
56  ();
57

```

Device

```
58 public ClockSlave: App ==> ClockSlave
59 ClockSlave(ap) ==
60 (
61   registerApp(ap);
62 );
63
64 public triggerPrint: () ==> ()
65 triggerPrint() ==
66   for all i in set {1,...,len timeConsumers} do
67     timeConsumers(i).showReceivedTime();
68
69 end ClockSlave
```

C.2.4 Device

```
1 class Device
2
3 instance variables
4
5   -- Device ID number
6   private deviceID : nat := 0;
7
8   -- Set device as a bridge
9   public isBridge : bool := true;
10
11   -- Frequency of the current device.
12   private freq : nat := 0;
13
14   -- Slave port associated to the device
15   public slavePort : Slave;
16
17   -- Device with no external slaves connected
18   -- (finalDevice = true)
19   public finalDevice : bool := false;
20   --inv finalDevice => len masterPorts = 0;
21
22   -- Current device is Grand Master -> iGM = true
23   -- if device isGM it should not have masters attached.
24   public isGM : bool := false;
25   inv isGM => not finalDevice;
26   inv finalDevice => not isGM;
27
28   -- Frequency of the master device
29   private masterFreq : nat := 0;
30
31   private receivedSyncEvent : bool := false;
```

Appendix C. AVB case study VDM models

```
32  private syncRecTS : nat := 0;
33
34  -- Neighbour rate ratio measured by the slave port
35  private neighbourRateRatio : real := 1;
36
37  -- Precise origin timestamp, defined by GM and
38  -- never changed.
39  private preciseOriginTS : real := 0;
40
41  -- received correction field, to be recomputed
42  -- before use.
43  private receivedCorrectionField : real := 0;
44
45  -- CumulativeRateRatio received from the previous
46  -- node.
47  private cummulativeRateRatio : real := 1;
48  -- Delay from master to slave port
49  private linkDelay : real := 10;
50  -- Time that has been spent in the current node
51  private residenceTime : real := 0;
52
53  -- Correction field calculated at the current node,
54  -- to be used in the current node.
55  private newCorrectionField : real := 0;
56  -- CummulativeRateRatio calculated at the current node,
57  -- to be used in the current node.
58  private newCummulativeRateRatio : real := 0;
59
60  -- Clock attached to the device
61  private clk : Clock;
62
63  private timeInfoAvailable : bool;
64
65  public cSlave : ClockSlave := new ClockSlave();
66
67  private followUpMessageReceived : bool := false;
68
69  operations
70
71  public setClockSlave: ClockSlave ==> ()
72  setClockSlave(c) ==
73    cSlave := c;
74
75  -- Actions carried out at the device per time-step
76  public timeStep: () ==> ()
77  timeStep() ==
78  (
79    if receivedSyncEvent = false then
80    (  receiveSyncEvent(); IO`print("#syncEvent#");)
```

Device

```
81  else if receivedSyncEvent and not followUpMessageReceived then
82    ( receiveFollowUpMessage(); IO`print("Follow up received");)
83  else if receivedSyncEvent and followUpMessageReceived then
84    ( notifyClock(); IO`print("CLocks notified"););
85  );
86
87
88  public notifyClock : () ==> ()
89  notifyClock() ==
90    cSlave.setTimeInfo(
91      [preciseOriginTS, receivedCorrectionField, syncRecTS]);
92
93  public makeTimeAvailable : () ==> ()
94  makeTimeAvailable() ==
95    timeInfoAvailable := true;
96
97  public getTimeInfo : () ==> seq of real
98  getTimeInfo() ==
99    return [
100      preciseOriginTS,
101      receivedCorrectionField,
102      cummulativeRateRatio];
103
104  public timeInfoRead : () ==> ()
105  timeInfoRead() ==
106    timeInfoAvailable := false;
107
108  -- Sets the slave port that corresponds to the current device
109  public setSlavePort: Slave ==> ()
110  setSlavePort(s) ==
111    slavePort := s;
112
113  -- Sets device ID
114  public setDeviceID: nat ==> ()
115  setDeviceID(id) ==
116    deviceID := id
117  pre deviceID = 0
118  post deviceID <> 0;
119
120  -- Gets device ID
121  public getDeviceID: () ==> nat
122  getDeviceID() ==
123    return deviceID
124  pre deviceID <> 0;
125
126  -- Receives the POTS
127  public receivePOTS: () ==> ()
128  receivePOTS() ==
129    preciseOriginTS := slavePort.receivePOTS();
```

Appendix C. AVB case study VDM models

```

130
131 public setLocalFreq : nat ==> ()
132 setLocalFreq(f) ==
133     freq := f;
134
135 -- Sets the frequency of the master device.
136 public setMasterFreq : nat ==> ()
137 setMasterFreq (f) ==
138     masterFreq := f
139 pre f <> 0;
140
141 -- Get the frequency of the master device attached to
142 -- the slave port.
143 public getFreq : () ==> ()
144 getFreq() ==
145     if (not isGM) then setMasterFreq(slavePort.getFreq());
146
147
148 -- Receives the follow up message sent by the master.
149 -- Reception is performed in one operation.
150 public receiveFollowUpMessage: () ==> ()
151 receiveFollowUpMessage() ==
152 ( let s : seq of real = slavePort.receiveFollowUpMessage() in
153   (
154     receivedCorrectionField := s(1);
155     preciseOriginTS := s(2);
156     cummulativeRateRatio := s(3);
157   );
158   followUpMessageReceived := true;
159 );
160
161 -- Receives the cummulative rate ratio at the slave
162 -- port. Cumulative rate ratio is sent by the master
163 -- device attached to the slave port.
164 public receiveCRR: () ==> ()
165 receiveCRR() ==
166     cummulativeRateRatio:= slavePort.receiveCRR();
167
168 -- Calculates the neighbour rate ratio. This calculation
169 -- is performed at every slave device.
170 public calculateNeighbourRR : () ==> ()
171 calculateNeighbourRR () ==
172     if (not isGM) then
173         neighbourRateRatio := freq / masterFreq
174     else
175         neighbourRateRatio := 1
176 pre not isGM => (freq <> 0 and masterFreq <> 0)
177 post neighbourRateRatio <> 0;
178

```

Device

```
179 -- Calculates the new cummulative rate ratio to be sent
180 -- to the attached slave unit.
181 public calculateCRateRatio : () ==> ()
182 calculateCRateRatio () ==
183     newCummulativeRateRatio :=
184         cummulativeRateRatio * neighbourRateRatio
185 pre neighbourRateRatio <> 0 and cummulativeRateRatio <> 0;
186
187 -- Calculates the new correction field to be sent to the
188 -- attached slave unit.
189 public calculateNewCorrectionField : () ==> ()
190 calculateNewCorrectionField () ==
191     if isGM then ()
192     else (
193         newCorrectionField := linkDelay + residenceTime *
194             newCummulativeRateRatio + receivedCorrectionField);
195
196 -- Receives a correction field sent by master unit.
197 public receiveCorrectionField : () ==> ()
198 receiveCorrectionField() ==
199     receivedCorrectionField := slavePort.receiveCorrectionField();
200
201 -- Sets the current device as a final device.
202 -- A final device will not have slave devices and
203 -- will not be acting as a master device.
204 public setFinalDevice: () ==> ()
205 setFinalDevice() ==
206     finalDevice := true
207 pre not isGM
208 post finalDevice;
209
210 -- Receives a synchronization sync event and
211 -- time stamps it.
212 public receiveSyncEvent: () ==> ()
213 receiveSyncEvent () ==
214 (
215     receivedSyncEvent := slavePort.receiveSyncEvent();
216     receivedSyncEvent := true;
217     updateTimeRef();
218 );
219
220 -- Updates the time reference when the time stamp operation
221 -- is invoked.
222 public updateTimeRef: () ==> ()
223 updateTimeRef() ==
224 (
225     --IO 'print("\n Sync event --> timestamp generated");
226     syncRecTS := slavePort.getIngressTime();
227 );
```

Appendix C. AVB case study VDM models

```
228 |
229 | -- Shows the device state
230 | public showDevice : () ==> ()
231 | showDevice () ==
232 | (
233 |   IO`print("\nInfo for device: ");
234 |   IO`print(deviceID);
235 |   IO`print("\nIs GM:");
236 |   IO`print(isGM);
237 |   IO`print("\nIs final device: ");
238 |   IO`print(finalDevice);
239 |   IO`print("\nIs sync event received: ");
240 |   IO`print(receivedSyncEvent);
241 |   IO`print("\nSync event received at: ");
242 |   IO`print(syncRecTS);
243 |   IO`print("\nFrequency: ");
244 |   IO`print(freq);
245 |   IO`print("\nMaster frequency: ");
246 |   IO`print(masterFreq);
247 |   IO`print("\nLinkDelay: ");
248 |   IO`print(linkDelay);
249 |
250 |   IO`print("\nNeighbour Rate Ratio: ");
251 |   IO`print(neighbourRateRatio);
252 |   IO`print("\nRecCummulative Rate Ratio: ");
253 |   IO`print(cummulativeRateRatio);
254 |
255 |   IO`print("\nReceived CorrectionField: ");
256 |   IO`print(receivedCorrectionField);
257 |   IO`print("\nPrecise Origin TimeStamp: ");
258 |   IO`print(preciseOriginTS);
259 | );
260 |
261 | end Device
```

C.2.5 Main

```
1 | class Main
2 |
3 | types
4 |
5 |   public timeStimulus = seq of real;
6 |
7 | instance variables
8 |
9 |   -- Modelled entities in the system
```


Main

```
10 | siteSync : Device;
11 | app : App;
12 | clkSlave : ClockSlave;
13 | slave : Slave;
14 | channel : PhyMessage;
15 |
16 | -- Auxiliar simulation entities
17 | simClk : Clock := new Clock();
18 |
19 | simSteps : nat := 3;
20 |
21 | timeStimuli : seq of timeStimulus := [];
22 |
23 | operations
24 |
25 | public loadStimuli: () ==> ()
26 | loadStimuli() ==
27 | (
28 |   --loadStimulus([20,30,50]);
29 |   let i : seq of timeStimulus = [[40,20,10],[30,80,90]] in
30 |     for all e in set {1,...,len i} do
31 |       loadStimulus(i(e));
32 | );
33 |
34 | public loadStimulus: timeStimulus ==> ()
35 | loadStimulus(s) ==
36 |   timeStimuli := timeStimuli ^ [s];
37 |
38 | -- Model simulation entry point
39 | public run: () ==> ()
40 | run() ==
41 | (
42 |   IO`print("\nEndpoint device receiving time corrections");
43 |
44 |   loadStimuli();
45 |   IO`print(timeStimuli);
46 |   setup();
47 |
48 |   -- Simulation starts
49 |
50 |   feedPhyMessage();
51 |
52 |   for all i in set {1,...,simSteps} do
53 |     (
54 |       simClk.tick();
55 |       slave.timeStep();
56 |       siteSync.timeStep();
57 |       clkSlave.timeStep();
58 |     );
```

Appendix C. AVB case study VDM models

```
59
60  -- Shows the simulation results
61  siteSync.showDevice();
62  app.showReceivedTime();
63
64  IO`print("\nModel over\n");
65  IO`print("##### Unused time stimuli: ");
66  IO`print(timeStimuli);
67
68 );
69
70 public feedPhyMessage: () ==> ()
71 feedPhyMessage() ==
72 (
73   channel.signalSyncEvent();
74   channel.setFollowUp(hd timeStimuli);
75   timeStimuli := tl timeStimuli;
76 );
77
78 -- Setup the simulation:
79 --   Creates the involved entities
80 --   Establish the relations between them
81 public setup: () ==> ()
82 setup() ==
83 (
84   -- Creation of the entities present in a Time Aware System
85   -- endpoint device
86   siteSync := new Device();
87   slave := new Slave();
88   clkSlave := new ClockSlave();
89   app := new App();
90   channel := new PhyMessage();
91
92   slave.clk := simClk;
93
94   -- Establish the relationships between the entities
95
96   -- 1. Register the Application in the clock slave entity
97   clkSlave.registerApp(app);
98   clkSlave.setSiteSync(siteSync);
99
100  -- 2. Associate the clock slave to the siteSync entity
101  siteSync.setClockSlave(clkSlave);
102
103  -- 3. Configure the physical layer and associate it to the
104  --   slave port
105  slave.changePhyChannel(channel);
106
107  -- 4. Associate the slave port to the siteSync entity
```

PhyMessage

```
108 | siteSync.setSlavePort(slave);
109 |
110 | channel.setMasterFreq(10);
111 | siteSync.setLocalFreq(20);
112 | siteSync.getFreq();
113 | siteSync.calculateNeighbourRR();
114 |
115 | siteSync.setFinalDevice();
116 | );
117 |
118 | end Main
```

C.2.6 PhyMessage

```
1 | class PhyMessage
2 |
3 | instance variables
4 |
5 |   -- sync event flag. Holds true value if a sync event
6 |   -- is put on the bus.
7 |   private syncEvent : bool := false;
8 |
9 |   -- Contents of the follow up message
10 |  private preciseOriginTS : real := 0.0;
11 |  private cummulativeRateRatio : real := 0.0;
12 |  private correctionField : real := 0.0;
13 |
14 |   -- Frequency sent by a master
15 |   private masterFreq : real := 0;
16 |
17 | operations
18 |
19 |   -- Gets the contents of the follow up message.
20 |  public getFollowUp: () ==> seq of real
21 |  getFollowUp() ==
22 |    return [correctionField,
23 |            preciseOriginTS,
24 |            cummulativeRateRatio];
25 |
26 |   -- Sets the contents of the follow up message.
27 |  public setFollowUp: seq of real ==> ()
28 |  setFollowUp(s) ==
29 |  (
30 |    setCRR(s(1));
31 |    setCorrectionField(s(2));
32 |    setPOTS(s(3));
```

Appendix C. AVB case study VDM models

```

33 );
34
35 -- Signals a Synchronization event
36 public signalSyncEvent: () ==> ()
37 signalSyncEvent() ==
38 (
39   --IO'print("\nMaster port --> Sync event sent");
40   syncEvent := true;
41 );
42
43 -- Check if there is a Synchronization event
44 public getSyncEvent: () ==> bool
45 getSyncEvent() ==
46 (
47   --IO'print("\nSlave port --> Sync event received");
48   return syncEvent;
49 );
50
51 -- #####
52 -- Getters and setters for the different fields in the
53 -- Physical Message.
54
55 public setMasterFreq : real ==> ()
56 setMasterFreq(f) ==
57   masterFreq := f;
58
59 public getMasterFreq: () ==> real
60 getMasterFreq() ==
61   return masterFreq;
62
63 public setCRR: real ==> ()
64 setCRR(crr) ==
65   cummulativeRateRatio := crr;
66
67 public setCorrectionField: real ==> ()
68 setCorrectionField(crf) ==
69   correctionField := crf;
70
71 public setPOTS: real ==> ()
72 setPOTS(pots) ==
73   preciseOriginTS := pots;
74
75 public getCField: () ==> real
76 getCField() ==
77   return correctionField;
78
79 public getPOTS: () ==> real
80 getPOTS() ==
81   return preciseOriginTS;

```

Port

```
82
83 public getCRR: () ==> real
84 getCRR() ==
85     return cummulativeRateRatio;
86
87 sync
88
89     mutex(getFollowUp, setFollowUp);
90     mutex(signalSyncEvent, getSyncEvent);
91
92 end PhyMessage
```

C.2.7 Port

```
1 class Port
2
3 instance variables
4
5     -- Physical channel associated to the port
6     protected channel : PhyMessage;
7
8     -- Associated ID to the port
9     protected portID : nat;
10
11     protected clk : Clock;
12
13 operations
14
15 public timeStamp : () ==> ()
16 timeStamp() ==
17     is not yet specified;
18
19 public timeStep : () ==> ()
20 timeStep() ==
21     is not yet specified;
22
23     -- Change the associated channel to the slave
24 public changePhyChannel : PhyMessage ==> ()
25 changePhyChannel(ch) ==
26     channel := ch;
27
28 public getPhyChannel : () ==> PhyMessage
29 getPhyChannel() ==
30     return channel;
31
32     -- Changes the associated ID to the port
```

```

33 public changePortID: nat ==> ()
34 changePortID(id) ==
35   portID := id;
36
37 end Port

```

C.2.8 Slave

```

1 class Slave is subclass of Port
2
3 instance variables
4
5   ingresTime : real := 0;
6
7   -- sync event flag. Holds true value if a sync event
8   -- is put on the bus.
9   private syncEvent : bool := false;
10
11   -- Contents of the follow up message
12   private preciseOriginTS : real := 0.0;
13   private cummulativeRateRatio : real := 0.0;
14   private correctionField : real := 0.0;
15
16   private followUpReady : bool := false;
17
18   public clk : Clock;
19
20
21 operations
22
23 public checkSync : () ==> ()
24 checkSync() ==
25 (
26   syncEvent := channel.getSyncEvent();
27   if syncEvent = true then
28     timeStamp(); -- Time stamp
29 );
30
31 public checkFollowUp: () ==> ()
32 checkFollowUp() ==
33 (
34   let i : seq of real = channel.getFollowUp() in
35   (
36     if len i = 0 then
37       followUpReady := false
38     else

```

Slave

```
39     (  
40         correctionField := i(1);  
41         preciseOriginTS := i(2);  
42         cummulativeRateRatio := i(3);  
43     )  
44 );  
45 );  
46  
47 public timeStamp : () ==> ()  
48 timeStamp() ==  
49     ingressTime := clk.getTime();  
50  
51 public getIngressTime : () ==> real  
52 getIngressTime () ==  
53     return ingressTime;  
54  
55 -- Class constructor  
56 public Slave: PhyMessage ==> Slave  
57 Slave(phy) ==  
58     channel := phy;  
59  
60 -- Receives the cummulative rate ratio  
61 public receiveCRR: () ==> real  
62 receiveCRR() ==  
63     channel.getCRR();  
64  
65 -- Receives the correction field  
66 public receiveCorrectionField: () ==> real  
67 receiveCorrectionField() ==  
68     channel.getCField();  
69  
70 -- Receives the precise origin time stamp  
71 public receivePOTS : () ==> real  
72 receivePOTS() ==  
73     channel.getPOTS();  
74  
75 -- Receives the follow up message  
76 -- Receives Precise Origing Time Stamp, cummulative rate ratio  
77 -- and correction field in one invocation  
78 public receiveFollowUpMessage: () ==> seq of real  
79 receiveFollowUpMessage() ==  
80     return [correctionField,preciseOriginTS,cummulativeRateRatio];  
81  
82 -- Checks if there is a synchronization event  
83 -- available  
84 public receiveSyncEvent: () ==> bool  
85 receiveSyncEvent() ==  
86     return syncEvent;  
87
```

```

88  -- Gets the frequency associated to the master
89  public getFreq: () ==> real
90  getFreq() ==
91    return channel.getMasterFreq();
92
93  public timeStep: () ==> ()
94  timeStep() ==
95  (
96    if not syncEvent then
97      checkSync()
98    else
99      if not followUpReady then
100        checkFollowUp();
101  );
102
103  end Slave

```

C.3. Concurrent model

C.3.1 App

```

1  class App
2
3  instance variables
4
5    private preciseOriginTS : real;
6    private correctionField : real;
7
8    private localTimeStamp : real;
9
10   private timeInfoReady : bool := false;
11
12   private timeInfoProvided : bool := false;
13
14
15  operations
16
17  public receiveTimeInformation : seq of real ==> ()
18  receiveTimeInformation (time) ==
19  (
20    preciseOriginTS := time(1);
21    correctionField := time(2);
22    localTimeStamp := time(3);
23

```


App

```
24   timeInfoReady := true;
25
26   markInfoProvided();
27
28 );
29
30 public showReceivedTime : () ==> ()
31 showReceivedTime() ==
32 (
33   if timeInfoReady then
34     (
35       IO`print("\n#####");
36       IO`print("\nAt time: ");
37       IO`print(localTimeStamp);
38       IO`print("\nThe Grand Master time was: ");
39       IO`print(preciseOriginTS);
40       IO`print("\nIt has been corrected with: ");
41       IO`print(correctionField);
42       IO`print("\nFinal time for previos timestamp is: ");
43       IO`print(preciseOriginTS+correctionField);
44       IO`print("\n#####");
45
46       markInfoProvided();
47     )
48   else
49     IO`print("\n Message from application:
50           Time information not Ready");
51 );
52
53 private markInfoProvided : () ==> ()
54 markInfoProvided() ==
55   timeInfoProvided := true;
56
57 public infoProvided : () ==> bool
58 infoProvided() ==
59   return timeInfoProvided;
60
61
62
63 sync
64   mutex(markInfoProvided,infoProvided);
65   mutex(receiveTimeInformation);
66   mutex(showReceivedTime, infoProvided);
67   mutex(showReceivedTime, receiveTimeInformation);
68
69 end App
```

C.3.2 Clock

```

1 class Clock
2
3 instance variables
4
5     time : nat := 0;
6
7 operations
8
9 public tick: () ==> ()
10 tick() ==
11 (
12     time := time +1;
13 );
14
15 public getTime: () ==> nat
16 getTime() ==
17     return time;
18
19 thread
20     while true do
21         (
22             tick();
23             Environment `timerRef.WaitRelative(1);
24         );
25
26 sync
27     mutex(tick,getTime);
28     mutex(tick);
29
30 end Clock

```

C.3.3 ClockSlave

```

1 class ClockSlave
2
3 instance variables
4
5     private siteSync : Device;
6
7     private preciseOriginTS : real;
8     private correctionField : real;
9     private localTimeStamp : real;
10

```

ClockSlave

```
11  private timeConsumers : seq of App := [];
12
13  private newTimeInfoReady : bool := false;
14
15  operations
16
17  public registerApp : App ==> ()
18  registerApp(ap) ==
19    timeConsumers := timeConsumers ^ [ap];
20
21  public setTimeInfo: seq of real ==> ()
22  setTimeInfo(time) ==
23  (
24    if len time > 2 then
25    (
26      preciseOriginTS := time(1);
27      correctionField := time(2);
28      localTimeStamp := time(3);
29      newTimeInfoReady := true;
30    )
31  );
32
33  public timeStep: () ==> ()
34  timeStep() ==
35    if newTimeInfoReady then
36    (
37      notifyTime();
38      newTimeInfoReady := false;
39    )
40    else
41      skip;
42
43  public notifyTime : () ==> ()
44  notifyTime() ==
45    let t : seq of real =
46      [preciseOriginTS, correctionField, localTimeStamp] in
47    for all i in set {1, ..., len timeConsumers} do
48      timeConsumers(i).receiveTimeInformation(t);
49
50  public setSiteSync : Device ==> ()
51  setSiteSync(dv) ==
52    siteSync := dv;
53
54  public getTime: () ==> ()
55  getTime() ==
56  ();
57
58  public ClockSlave: App ==> ClockSlave
59  ClockSlave(ap) ==
```

Appendix C. AVB case study VDM models

```
60 (
61   registerApp(ap);
62 );
63
64 public triggerPrint: () ==> ()
65 triggerPrint() ==
66   for all i in set {1,...,len timeConsumers} do
67     timeConsumers(i).showReceivedTime();
68
69 thread
70   while true do
71     (
72       timeStep();
73       Environment `timerRef.WaitRelative(1);
74     );
75
76 sync
77   mutex(timeStep, setTimeInfo);
78   mutex(notifyTime, setTimeInfo);
79
80 end ClockSlave
```

C.3.4 Device

```
1 class Device
2
3 instance variables
4
5   -- Device ID number
6   private deviceID : nat := 0;
7
8   -- Set device as a bridge
9   public isBridge : bool := true;
10
11   -- Frequency of the current device.
12   private freq : nat := 0;
13
14   -- Slave port associated to the device
15   public slavePort : Slave;
16
17   -- Device with no external slaves connected
18   -- (finalDevice = true)
19   public finalDevice : bool := false;
20   --inv finalDevice => len masterPorts = 0;
21
22   -- Current device is Grand Master -> iGM = true
```

Device

```
23  -- if device isGM it should not have masters attached.
24  public isGM : bool := false;
25  inv isGM => not finalDevice;
26  inv finalDevice => not isGM;
27
28  -- Frequency of the master device
29  private masterFreq : nat := 0;
30
31  private receivedSyncEvent : bool := false;
32  private syncRecTS : nat := 0;
33
34  -- Neighbour rate ratio measured by the slave port
35  private neighbourRateRatio : real := 1;
36
37  -- Precise origin timestamp, defined by GM and
38  -- never changed.
39  private preciseOriginTS : real := 0;
40
41  -- received correction field, to be recomputed
42  -- before use.
43  private receivedCorrectionField : real := 0;
44
45  -- CumulativeRateRatio received from the previous
46  -- node.
47  private cummulativeRateRatio : real := 1;
48  -- Delay from master to slave port
49  private linkDelay : real := 10;
50  -- Time that has been spent in the current node
51  private residenceTime : real := 0;
52
53  -- Correction field calculated at the current node,
54  -- to be used in the current node.
55  private newCorrectionField : real := 0;
56  -- CumulativeRateRatio calculated at the current node,
57  -- to be used in the current node.
58  private newCummulativeRateRatio : real := 0;
59
60  -- Clock attached to the device
61  private clk : Clock;
62
63  private timeInfoAvailable : bool;
64
65  public cSlave : ClockSlave := new ClockSlave();
66
67  private followUpMessageReceived : bool := false;
68
69  operations
70
71  public setClockSlave: ClockSlave ==> ()
```

Appendix C. AVB case study VDM models

```

72 setClockSlave(c) ==
73     cSlave := c;
74
75 -- Actions carried out at the device per time-step
76 public timeStep: () ==> ()
77 timeStep() ==
78 (
79     if receivedSyncEvent = false then
80     ( receiveSyncEvent(); IO`print("#syncEvent#");)
81     else if receivedSyncEvent and not followUpMessageReceived then
82     ( receiveFollowUpMessage(); IO`print("Follow up received");)
83     else if receivedSyncEvent and followUpMessageReceived then
84     ( notifyClock(); IO`print("CLocks notified"););
85 );
86
87
88 public notifyClock : () ==> ()
89 notifyClock() ==
90     cSlave.setTimeInfo(
91         [preciseOriginTS, receivedCorrectionField, syncRecTS]
92     );
93
94 public makeTimeAvailable : () ==> ()
95 makeTimeAvailable() ==
96     timeInfoAvailable := true;
97
98 public getTimeInfo : () ==> seq of real
99 getTimeInfo() ==
100     return [preciseOriginTS,
101             receivedCorrectionField,
102             cummulativeRateRatio];
103
104 public timeInfoRead : () ==> ()
105 timeInfoRead() ==
106     timeInfoAvailable := false;
107
108 -- Sets the slave port that corresponds to the current device
109 public setSlavePort: Slave ==> ()
110 setSlavePort(s) ==
111     slavePort := s;
112
113 -- Sets device ID
114 public setDeviceID: nat ==> ()
115 setDeviceID(id) ==
116     deviceID := id
117 pre deviceID = 0
118 post deviceID <> 0;
119
120 -- Gets device ID

```

Device

```
121 public getDeviceID: () ==> nat
122 getDeviceID() ==
123     return deviceID
124 pre deviceID <> 0;
125
126 -- Receives the POTS
127 public receivePOTS: () ==> ()
128 receivePOTS() ==
129     preciseOriginTS := slavePort.receivePOTS();
130
131 public setLocalFreq : nat ==> ()
132 setLocalFreq(f) ==
133     freq := f;
134
135 -- Sets the frequency of the master device.
136 public setMasterFreq : nat ==> ()
137 setMasterFreq (f) ==
138     masterFreq := f
139 pre f <> 0;
140
141 -- Get the frequency of the master device attached to
142 -- the slave port.
143 public getFreq : () ==> ()
144 getFreq() ==
145     if (not isGM) then setMasterFreq(slavePort.getFreq());
146
147
148 -- Receives the follow up message sent by the master.
149 -- Reception is performed in one operation.
150 public receiveFollowUpMessage: () ==> ()
151 receiveFollowUpMessage() ==
152 ( let s : seq of real = slavePort.receiveFollowUpMessage() in
153   (
154     receivedCorrectionField := s(1);
155     preciseOriginTS := s(2);
156     cummulativeRateRatio := s(3);
157   );
158   followUpMessageReceived := true;
159 );
160
161 -- Receives the cummulative rate ratio at the slave
162 -- port. Cummulative rate ratio is sent by the master
163 -- device attached to the slave port.
164 public receiveCRR: () ==> ()
165 receiveCRR() ==
166     cummulativeRateRatio:= slavePort.receiveCRR();
167
168 -- Calculates the neighbour rate ratio. This calculation
169 -- is performed at every slave device.
```

Appendix C. AVB case study VDM models

```

170 public calculateNeighbourRR : () ==> ()
171 calculateNeighbourRR () ==
172   if (not isGM) then
173     neighbourRateRatio := freq / masterFreq
174   else
175     neighbourRateRatio := 1
176 pre not isGM => (freq <> 0 and masterFreq <> 0)
177 post neighbourRateRatio <> 0;
178
179 -- Calculates the new cumulative rate ratio to be sent
180 -- to the attached slave unit.
181 public calculateCRateRatio : () ==> ()
182 calculateCRateRatio () ==
183   newCummulativeRateRatio :=
184     cummulativeRateRatio * neighbourRateRatio
185 pre neighbourRateRatio <> 0 and cummulativeRateRatio <> 0;
186
187 -- Calculates the new correction field to be sent to the
188 -- attached slave unit.
189 public calculateNewCorrectionField : () ==> ()
190 calculateNewCorrectionField () ==
191   if isGM then ()
192   else (
193     newCorrectionField := linkDelay + residenceTime *
194     newCummulativeRateRatio + receivedCorrectionField;);
195
196 -- Receives a correction field sent by master unit.
197 public receiveCorrectionField : () ==> ()
198 receiveCorrectionField() ==
199   receivedCorrectionField := slavePort.receiveCorrectionField();
200
201 -- Sets the current device as a final device.
202 -- A final device will not have slave devices and
203 -- will not be acting as a master device.
204 public setFinalDevice: () ==> ()
205 setFinalDevice() ==
206   finalDevice := true
207 pre not isGM
208 post finalDevice;
209
210 -- Receives a synchronization sync event and
211 -- time stamps it.
212 public receiveSyncEvent: () ==> ()
213 receiveSyncEvent () ==
214 (
215   receivedSyncEvent := slavePort.receiveSyncEvent();
216   receivedSyncEvent := true;
217   updateTimeRef();
218 );

```


Device

```
219
220 -- Updates the time reference when the time stamp operation
221 -- is invoked.
222 public updateTimeRef: () ==> ()
223 updateTimeRef() ==
224 (
225     --IO`print("\n Sync event --> timestamp generated");
226     syncRecTS := slavePort.getIngressTime();
227 );
228
229 -- Shows the device state
230 public showDevice : () ==> ()
231 showDevice () ==
232 (
233     IO`print("\nInfo for device: ");
234     IO`print(deviceID);
235     IO`print("\nIs GM:");
236     IO`print(isGM);
237     IO`print("\nIs final device: ");
238     IO`print(finalDevice);
239     IO`print("\nIs sync event received: ");
240     IO`print(receivedSyncEvent);
241     IO`print("\nSync event received at: ");
242     IO`print(syncRecTS);
243     IO`print("\nFrequency: ");
244     IO`print(freq);
245     IO`print("\nMaster frequency: ");
246     IO`print(masterFreq);
247     IO`print("\nLinkDelay: ");
248     IO`print(linkDelay);
249
250     IO`print("\nNeighbour Rate Ratio: ");
251     IO`print(neighbourRateRatio);
252     IO`print("\nRecCummulative Rate Ratio: ");
253     IO`print(cummulativeRateRatio);
254
255     IO`print("\nReceived CorrectionField: ");
256     IO`print(receivedCorrectionField);
257     IO`print("\nPrecise Origin TimeStamp: ");
258     IO`print(preciseOriginTS);
259
260 );
261
262 thread
263     while true do
264     (
265         timeStep();
266         Environment`timerRef.WaitRelative(1);
267     );
```

268
269 **end** Device

C.3.5 Environment

```

1 class Environment
2
3 types
4
5   public timeStimulus = seq of real;
6
7 instance variables
8
9   -- Modelled entities in the system
10  siteSync : Device;
11  app : App;
12  clkSlave : ClockSlave;
13  slave : Slave;
14  channel : PhyMessage;
15
16  -- Auxiliar simulation entities
17  simClk : Clock := new Clock();
18
19  simSteps : nat := 3;
20
21  timeStimuli : seq of timeStimulus := [];
22
23  public static timerRef : TimeStamp := new TimeStamp(5);
24
25 operations
26
27 public loadStimuli: () ==> ()
28 loadStimuli() ==
29 (
30   --loadStimulus([20,30,50]);
31   let i : seq of timeStimulus = [[40,20,10],[30,80,90]] in
32     for all e in set {1,...,len i} do
33       loadStimulus(i(e));
34 );
35
36 public loadStimulus: timeStimulus ==> ()
37 loadStimulus(s) ==
38   timeStimuli := timeStimuli ^ [s];
39
40 -- Model simulation entry point
41 public run: () ==> ()

```

Environment

```
42 run() ==
43 (
44   IO`print("\nEndpoint device receiving time corrections");
45
46   loadStimuli();
47   IO`print(timeStimuli);
48   setup();
49
50   -- Simulation starts
51
52   start (simClk);
53   start (slave);
54   start (siteSync);
55   start (clkSlave);
56
57   feedPhyMessage();
58
59   /*
60     for all i in set {1,...,simSteps} do
61     (
62       simClk.tick();
63       slave.timeStep();
64       siteSync.timeStep();
65       clkSlave.timeStep();
66     );*/
67
68     while app.infoProvided() = false do
69     (
70       wait();
71       timerRef.WaitRelative(1);
72     );
73
74     -- Shows the simulation results
75     siteSync.showDevice();
76     app.showReceivedTime();
77
78     IO`print("\nModel over\n");
79     IO`print("##### Unused time stimuli: ");
80     IO`print(timeStimuli);
81
82 );
83
84 public feedPhyMessage: () ==> ()
85 feedPhyMessage() ==
86 (
87   channel.signalSyncEvent();
88   channel.setFollowUp(hd timeStimuli);
89   timeStimuli := tl timeStimuli;
90 );
```

Appendix C. AVB case study VDM models

```
91 |
92 | -- Setup the simulation:
93 | --   Creates the involved entities
94 | --   Establish the relations between them
95 | public setup: () ==> ()
96 | setup() ==
97 | (
98 |   -- Creation of the entities present in a Time Aware System
99 |   -- endpoint device
100 |   siteSync := new Device();
101 |   slave := new Slave();
102 |   clkSlave := new ClockSlave();
103 |   app := new App();
104 |   channel := new PhyMessage();
105 |
106 |   slave.clk := simClk;
107 |
108 |   -- Establish the relationships between the entities
109 |
110 |   -- 1. Register the Application in the clock slave entity
111 |   clkSlave.registerApp(app);
112 |   clkSlave.setSiteSync(siteSync);
113 |
114 |   -- 2. Associate the clock slave to the siteSync entity
115 |   siteSync.setClockSlave(clkSlave);
116 |
117 |   -- 3. Configure the physical layer and associate it to the
118 |   --   slave port
119 |   slave.changePhyChannel(channel);
120 |
121 |   -- 4. Associate the slave port to the siteSync entity
122 |   siteSync.setSlavePort(slave);
123 |
124 |   channel.setMasterFreq(10);
125 |   siteSync.setLocalFreq(20);
126 |   siteSync.getFreq();
127 |   siteSync.calculateNeighbourRR();
128 |
129 |   siteSync.setFinalDevice();
130 | );
131 |
132 | wait: () ==> ()
133 | wait() ==
134 |   skip;
135 |
136 | end Environment
```

C.3.6 PhyMessage

```

1 class PhyMessage
2
3 instance variables
4
5   -- sync event flag. Holds true value if a sync event
6   -- is put on the bus.
7   private syncEvent : bool := false;
8
9   -- Contents of the follow up message
10  private preciseOriginTS : real := 0.0;
11  private cumulativeRateRatio : real := 0.0;
12  private correctionField : real := 0.0;
13
14  -- Frequency sent by a master
15  private masterFreq : real := 0;
16
17 operations
18
19  -- Gets the contents of the follow up message.
20  public getFollowUp: () ==> seq of real
21  getFollowUp() ==
22    return [correctionField,
23            preciseOriginTS,
24            cumulativeRateRatio];
25
26  -- Sets the contents of the follow up message.
27  public setFollowUp: seq of real ==> ()
28  setFollowUp(s) ==
29  (
30    setCRR(s(1));
31    setCorrectionField(s(2));
32    setPOTS(s(3));
33  );
34
35  -- Signals a Synchronization event
36  public signalSyncEvent: () ==> ()
37  signalSyncEvent() ==
38  (
39    --IO`print("\nMaster port --> Sync event sent");
40    syncEvent := true;
41  );
42
43  -- Check if there is a Synchronization event
44  public getSyncEvent: () ==> bool
45  getSyncEvent() ==
46  (

```

Appendix C. AVB case study VDM models

```

47  --IO`print("\nSlave port --> Sync event received");
48  return syncEvent;
49  );
50
51  -- #####
52  -- Getters and setters for the different fields in the
53  -- Physical Message.
54
55  public setMasterFreq : real ==> ()
56  setMasterFreq(f) ==
57    masterFreq := f;
58
59  public getMasterFreq: () ==> real
60  getMasterFreq() ==
61    return masterFreq;
62
63  public setCRR: real ==> ()
64  setCRR(crr) ==
65    cumulativeRateRatio := crr;
66
67  public setCorrectionField: real ==> ()
68  setCorrectionField(crf) ==
69    correctionField := crf;
70
71  public setPOTS: real ==> ()
72  setPOTS(pots) ==
73    preciseOriginTS := pots;
74
75  public getCField: () ==> real
76  getCField() ==
77    return correctionField;
78
79  public getPOTS: () ==> real
80  getPOTS() ==
81    return preciseOriginTS;
82
83  public getCRR: () ==> real
84  getCRR() ==
85    return cumulativeRateRatio;
86
87  sync
88
89    mutex(getFollowUp, setFollowUp);
90    mutex(signalSyncEvent, getSyncEvent);
91
92  end PhyMessage

```

C.3.7 Port

```

1 class Port
2
3 instance variables
4
5   -- Physical channel associated to the port
6   protected channel : PhyMessage;
7
8   -- Associated ID to the port
9   protected portID : nat;
10
11  protected clk : Clock;
12
13 operations
14
15 public timeStamp : () ==> ()
16 timeStamp() ==
17   is not yet specified;
18
19 public timeStep : () ==> ()
20 timeStep() ==
21   is not yet specified;
22
23   -- Change the associated channel to the slave
24 public changePhyChannel : PhyMessage ==> ()
25 changePhyChannel(ch) ==
26   channel := ch;
27
28 public getPhyChannel : () ==> PhyMessage
29 getPhyChannel() ==
30   return channel;
31
32   -- Changes the associated ID to the port
33 public changePortID: nat ==> ()
34 changePortID(id) ==
35   portID := id;
36
37 end Port

```

C.3.8 Slave

```

1 class Slave is subclass of Port
2
3 instance variables

```

Appendix C. AVB case study VDM models

```

4
5  ingresTime : real := 0;
6
7  -- sync event flag. Holds true value if a sync event
8  -- is put on the bus.
9  private syncEvent : bool := false;
10
11  -- Contents of the follow up message
12  private preciseOriginTS : real := 0.0;
13  private cummulativeRateRatio : real := 0.0;
14  private correctionField : real := 0.0;
15
16  private followUpReady : bool := false;
17
18  public clk : Clock;
19
20
21 operations
22
23 public checkSync : () ==> ()
24 checkSync() ==
25 (
26   syncEvent := channel.getSyncEvent();
27   if syncEvent = true then
28     timeStamp(); -- Time stamp
29 );
30
31 public checkFollowUp: () ==> ()
32 checkFollowUp() ==
33 (
34   let i : seq of real = channel.getFollowUp() in
35   (
36     if len i = 0 then
37       followUpReady := false
38     else
39       (
40         correctionField := i(1);
41         preciseOriginTS := i(2);
42         cummulativeRateRatio := i(3);
43       )
44   );
45 );
46
47 public timeStamp : () ==> ()
48 timeStamp() ==
49   ingresTime := clk.getTime();
50
51 public getIngresTime : () ==> real
52 getIngresTime () ==

```


Slave

```
53   return ingressTime;
54
55   -- Class constructor
56   public Slave: PhyMessage ==> Slave
57   Slave(phy) ==
58     channel := phy;
59
60   -- Receives the cummulative rate ratio
61   public receiveCRR: () ==> real
62   receiveCRR() ==
63     channel.getCRR();
64
65   -- Receives the correction field
66   public receiveCorrectionField: () ==> real
67   receiveCorrectionField() ==
68     channel.getCField();
69
70   -- Receives the precise origin time stamp
71   public receivePOTS : () ==> real
72   receivePOTS() ==
73     channel.getPOTS();
74
75   -- Receives the follow up message
76   -- Receives Precise Origing Time Stamp, cummulative rate ratio
77   -- and correction field in one invocation
78   public receiveFollowUpMessage: () ==> seq of real
79   receiveFollowUpMessage() ==
80     return [correctionField,preciseOriginTS,cummulativeRateRatio];
81
82   -- Checks if there is a synchronization event
83   -- available
84   public receiveSyncEvent: () ==> bool
85   receiveSyncEvent() ==
86     return syncEvent;
87
88   -- Gets the frequency associated to the master
89   public getFreq: () ==> real
90   getFreq() ==
91     return channel.getMasterFreq();
92
93   public timeStep: () ==> ()
94   timeStep() ==
95   (
96     if not syncEvent then
97       checkSync()
98     else
99       if not followUpReady then
100         checkFollowUp();
101 );
```

```

102
103 thread
104   while true do
105     (
106       timeStep();
107       Environment `timerRef.WaitRelative(1);
108     );
109
110 end Slave

```

C.3.9 TimeStamp

```

1 class TimeStamp
2
3 values
4
5 public stepLength : nat = 1;
6
7 instance variables
8
9 currentTime : nat := 0;
10 wakeUpMap : map nat to [nat] := {|->};
11 barrierCount : nat1;
12
13 operations
14
15 public TimeStamp : nat1 ==> TimeStamp
16 TimeStamp(count) ==
17   barrierCount := count;
18
19 public WaitRelative : nat ==> ()
20 WaitRelative(val) ==
21   WaitAbsolute(currentTime + val);
22
23 public WaitAbsolute : nat ==> ()
24 WaitAbsolute(val) == (
25   AddToWakeUpMap(threadid, val);
26   -- Last to enter the barrier notifies the rest.
27   BarrierReached();
28   -- Wait till time is up
29   Awake();
30 );
31
32 BarrierReached : () ==> ()
33 BarrierReached() ==
34 (

```

TimeStamp

```

35  while (card dom wakeUpMap = barrierCount) do
36    (
37      currentTime := currentTime + stepLength;
38      let threadSet : set of nat = {th | th in set dom wakeUpMap
39        & wakeUpMap(th) <> nil and
40        wakeUpMap(th) <= currentTime }
41      in
42        for all t in set threadSet
43        do
44          wakeUpMap := {t} <-: wakeUpMap;
45        );
46    )
47  post forall x in set rng wakeUpMap & x = nil or
48    x >= currentTime;
49
50  AddToWakeUpMap : nat * [nat] ==> ()
51  AddToWakeUpMap(tId, val) ==
52    wakeUpMap := wakeUpMap ++ { tId |-> val };
53
54  public NotifyThread : nat ==> ()
55  NotifyThread(tId) ==
56    wakeUpMap := {tId} <-: wakeUpMap;
57
58  public GetTime : () ==> nat
59  GetTime() ==
60    return currentTime;
61
62  Awake: () ==> ()
63  Awake() == skip;
64
65  public ThreadDone : () ==> ()
66  ThreadDone() ==
67    AddToWakeUpMap(threadid, nil);
68
69  sync
70    per Awake => threadid not in set dom wakeUpMap;
71
72    mutex(AddToWakeUpMap);
73    mutex(NotifyThread);
74    mutex(BarrierReached);
75
76    mutex(AddToWakeUpMap, NotifyThread);
77    mutex(AddToWakeUpMap, BarrierReached);
78    mutex(NotifyThread, BarrierReached);
79
80    mutex(AddToWakeUpMap, NotifyThread, BarrierReached);
81
82  end TimeStamp

```

C.4. Real-time model

C.4.1 App

```

1  class App
2
3  instance variables
4
5     private preciseOriginTS : real;
6     private correctionField : real;
7
8     private localTimeStamp : real;
9
10    private timeInfoReady : bool := false;
11
12    private timeInfoProvided : bool := false;
13
14
15  operations
16
17  public receiveTimeInformation : seq of real ==> ()
18  receiveTimeInformation (timeSeq) ==
19  (
20     preciseOriginTS := timeSeq(1);
21     correctionField := timeSeq(2);
22     localTimeStamp := timeSeq(3);
23
24     timeInfoReady := true;
25
26     markInfoProvided();
27
28  );
29
30  public showReceivedTime : () ==> ()
31  showReceivedTime () ==
32  (
33     if timeInfoReady then
34     (
35     IO`print("\nAt time: ");
36     IO`print(localTimeStamp);
37     IO`print("\nThe Grand Master time was: ");
38     IO`print(preciseOriginTS);
39     IO`print("\nIt has been corrected with: ");
40     IO`print(correctionField);
41     IO`print("\nFinal time for previos timestamp is: ");
42     IO`print(preciseOriginTS+correctionField);
43
44     markInfoProvided();

```

Clock

```
45 )
46 else
47   IO`print("\n Message from application:
48     Time information not Ready");
49 );
50
51 private markInfoProvided : () ==> ()
52 markInfoProvided() ==
53   timeInfoProvided := true;
54
55 public infoProvided : () ==> bool
56 infoProvided() ==
57   return timeInfoProvided;
58
59
60
61 sync
62   mutex(markInfoProvided, infoProvided);
63   mutex(receiveTimeInformation);
64   mutex(showReceivedTime, infoProvided);
65   mutex(showReceivedTime, receiveTimeInformation);
66
67 end App
```

C.4.2 Clock

```
1 class Clock
2
3 instance variables
4
5   currentTime : nat := 0;
6
7 operations
8
9 public tick: () ==> ()
10 tick() ==
11 (
12   duration (0) currentTime := currentTime +1;
13 );
14
15 public getTime: () ==> nat
16 getTime() ==
17   return time;
18
19 thread
20   -- while true do
```

Appendix C. AVB case study VDM models

```
21  -- (  
22  --  tick();  
23  -- );  
24  
25  -- periodic (2000E6,0,0,0) (tick)  
26  -- periodic (9E9,0,0,0) (tick) (Working)  
27  periodic(1E1,0,0,0) (tick)  
28  
29  -- periodic (1E3,0,0,0) (tick)  
30  
31  sync  
32  mutex(tick,getTime);  
33  mutex(tick);  
34  
35  end Clock
```

C.4.3 ClockSlave

```
1  class ClockSlave  
2  
3  instance variables  
4  
5  private siteSync : Device;  
6  
7  private preciseOriginTS : real;  
8  private correctionField : real;  
9  private localTimeStamp : real;  
10  
11  private timeConsumers : seq of App := [];  
12  
13  private newTimeInfoReady : bool := false;  
14  
15  operations  
16  
17  public registerApp : App ==> ()  
18  registerApp(ap) ==  
19  timeConsumers := timeConsumers ^ [ap];  
20  
21  public setTimeInfo: seq of real ==> ()  
22  setTimeInfo(timeSeq) ==  
23  (  
24  if len timeSeq > 2 then  
25  (  
26  preciseOriginTS := timeSeq(1);  
27  correctionField := timeSeq(2);  
28  localTimeStamp := timeSeq(3);
```

ClockSlave

```
29     newTimeInfoReady := true;
30   )
31
32 );
33
34 public timeStep: () ==> ()
35 timeStep() ==
36   if newTimeInfoReady then
37     (
38       notifyTime();
39       newTimeInfoReady := false;
40     )
41   else
42     skip;
43
44 public notifyTime : () ==> ()
45 notifyTime() ==
46   let t : seq of real =
47     [preciseOriginTS,
48       correctionField,
49       localTimeStamp] in
50   for all i in set {1,...,len timeConsumers} do
51     timeConsumers(i).receiveTimeInformation(t);
52
53 public setSiteSync : Device ==> ()
54 setSiteSync(dv) ==
55   siteSync := dv;
56
57 public getTime:() ==> ()
58 getTime() ==
59   ();
60
61 public ClockSlave: App ==> ClockSlave
62 ClockSlave(ap) ==
63   (
64     registerApp(ap);
65   );
66
67 public triggerPrint: () ==> ()
68 triggerPrint() ==
69   for all i in set {1,...,len timeConsumers} do
70     timeConsumers(i).showReceivedTime();
71
72 thread
73   while true do
74     (
75       timeStep();
76     );
77
```

```

78 sync
79   mutex(timeStep, setTimeInfo);
80   mutex(notifyTime, setTimeInfo);
81
82 end ClockSlave

```

C.4.4 Deployer

```

1  system Deployer
2
3  instance variables
4
5  -- Architecture definition
6
7  /*
8  Computing units:
9  Priority: <FP> - Fixed priority
10           <PP> - Priority?
11
12 Speed is giving in MIPS - Millions of instructions per second
13 */
14
15   cpu1 : CPU := new CPU(<FP>, 1E7);
16   cpu2 : CPU := new CPU(<FP>, 1E9);
17
18   cpu3 : CPU := new CPU(<FP>, 1E7);
19
20   cpu4 : CPU := new CPU(<FP>, 1E9);
21
22 /*
23 Communication bus:
24 Modes: <CSMACD> - ?
25 */
26   bus : BUS := new BUS(<CSMACD>, 72E9, { cpu1,cpu2});
27   bus2 : BUS := new BUS(<CSMACD>, 20E5, {cpu1,cpu3});
28
29
30   -- Used only to evaluate the fourth configuration
31   --(independent SiteSync block)
32   --bus : BUS := new BUS(<CSMACD>, 72E9, { cpu1,cpu4});
33
34   siteSyncCom : BUS := new BUS(<CSMACD>, 72E9, {cpu4,cpu2});
35
36
37   public static siteSync : Device := new Device();
38   public static slave : Slave := new Slave();

```


Deployer

```
39  public static clkSlave : ClockSlave := new ClockSlave();
40  public static app : App := new App();
41  public static channel : PhyMessage := new PhyMessage();
42  public static simClk : Clock := new Clock();
43  public static timerRef : TimeStamp := new TimeStamp(1);
44
45  operations
46
47  public Deployer : () ==> Deployer
48  Deployer () ==
49  (-- TODO Deploy deployable object to cpu's
50
51   slave.clk := simClk;
52
53   -- Establish the relationships between the entities
54
55   -- 1. Register the Application in the clock slave entity
56   clkSlave.registerApp(app);
57   clkSlave.setSiteSync(siteSync);
58
59   -- 2. Associate the clock slave to the siteSync entity
60   siteSync.setClockSlave(clkSlave);
61
62   -- 3. Configure the physical layer and associate it to the
63   slave port
64   slave.changePhyChannel(channel);
65
66   -- 4. Associate the slave port to the siteSync entity
67   siteSync.setSlavePort(slave);
68
69   channel.setMasterFreq(10);
70   siteSync.setLocalFreq(20);
71   siteSync.getFreq();
72   siteSync.calculateNeighbourRR();
73
74   siteSync.setFinalDevice();
75
76   -- Configuration 1
77   cpu1.deploy(slave);
78   cpu1.deploy(clkSlave);
79   cpu1.deploy(app);
80   cpu1.deploy(siteSync);
81
82   cpu2.deploy(channel);
83   cpu2.deploy(simClk);
84
85
86  /*
87  -- Architecture 2
```

Appendix C. AVB case study VDM models

```
88   cpu1.deploy(app);
89   cpu1.deploy(clkSlave);
90   cpu1.deploy(siteSync);
91
92   cpu2.deploy(slave);
93   cpu2.deploy(channel);
94   cpu2.deploy(simClk);
95 */
96
97 -- Architecture 3
98 /*
99   cpu1.deploy(clkSlave);
100  cpu1.deploy(app);
101
102  cpu2.deploy(siteSync);
103  cpu2.deploy(slave);
104  cpu2.deploy(channel);
105  cpu2.deploy(simClk);
106 */
107
108 -- Architecture 4
109
110 /*
111   cpu1.deploy(clkSlave);
112   cpu1.deploy(app);
113
114   cpu4.deploy(siteSync);
115
116   cpu2.deploy(slave);
117   cpu2.deploy(channel);
118   cpu2.deploy(simClk);
119 */
120
121 -- CPU's are started implicit
122 );
123
124 end Deployer
```

C.4.5 Device

```
1 class Device
2
3 instance variables
4
5   -- Device ID number
6   private deviceID : nat := 0;
```

Device

```
7
8  -- Set device as a bridge
9  public isBridge : bool := true;
10
11  -- Frequency of the current device.
12  private freq : nat := 0;
13
14  -- Slave port associated to the device
15  public slavePort : Slave;
16
17  -- Device with no external slaves connected
18  -- (finalDevice = true)
19  public finalDevice : bool := false;
20  --inv finalDevice => len masterPorts = 0;
21
22  -- Current device is Grand Master -> iGM = true
23  -- if device isGM it should not have masters attached.
24  public isGM : bool := false;
25  inv isGM => not finalDevice;
26  inv finalDevice => not isGM;
27
28  -- Frequency of the master device
29  private masterFreq : nat := 0;
30
31  private receivedSyncEvent : bool := false;
32  private syncRecTS : nat := 0;
33
34  -- Neighbour rate ratio measured by the slave port
35  private neighbourRateRatio : real := 1;
36
37  -- Precise origin timestamp, defined by GM and
38  -- never changed.
39  private preciseOriginTS : real := 0;
40
41  -- received correction field, to be recomputed
42  -- before use.
43  private receivedCorrectionField : real := 0;
44
45  -- CumulativeRateRatio received from the previous
46  -- node.
47  private cummulativeRateRatio : real := 1;
48  -- Delay from master to slave port
49  private linkDelay : real := 10;
50  -- Time that has been spent in the current node
51  private residenceTime : real := 0;
52
53  -- Correction field calculated at the current node,
54  -- to be used in the current node.
55  private newCorrectionField : real := 0;
```

Appendix C. AVB case study VDM models

```

56  -- CumulativeRateRatio calculated at the current node,
57  -- to be used in the current node.
58  private newCumulativeRateRatio : real := 0;
59
60  -- Clock attached to the device
61  private clk : Clock;
62
63  private timeInfoAvailable : bool;
64
65  public cSlave : ClockSlave := new ClockSlave();
66
67  private followUpMessageReceived : bool := false;
68
69  private clockNotified : bool := false;
70
71
72  private readSyncEvent : bool := false;
73  private readFollowUpMessage : bool := false;
74
75  operations
76
77  public setClockSlave: ClockSlave ==> ()
78  setClockSlave(c) ==
79    cSlave := c;
80
81  -- Actions carried out at the device per time-step
82  public timeStep: () ==> ()
83  timeStep() ==
84  (
85    if receivedSyncEvent = false then
86      ( receiveSyncEvent(); IO`print("#syncEvent#");)
87    else if receivedSyncEvent and not followUpMessageReceived then
88      ( receiveFollowUpMessage(); IO`print("Follow up received");)
89    else if receivedSyncEvent and followUpMessageReceived then
90      ( notifyClock(); IO`print("CLocks notified"););
91  );
92
93
94  -- Actions carried out at the device per time-step
95  public timeStep2: () ==> ()
96  timeStep2() ==
97  (
98    if receivedSyncEvent = false then
99      ( receiveSyncEvent();)
100    else if receivedSyncEvent and
101      not followUpMessageReceived then
102      ( receiveFollowUpMessage(); )
103    else if receivedSyncEvent and
104      followUpMessageReceived and

```

```

105     not clockNotified then
106     (  notifyClock(); );
107 );
108
109 public markReadSE : () ==> ()
110 markReadSE() ==
111     readSyncEvent := true;
112
113 public markReadFUM : () ==> ()
114 markReadFUM() ==
115     readFollowUpMessage := true;
116
117 public readyForNext : () ==> ()
118 readyForNext() ==
119 (
120     readFollowUpMessage := false;
121     readSyncEvent := false;
122 );
123
124 public notifyClock : () ==> ()
125 notifyClock() ==
126 (
127     cSlave.setTimeInfo(
128         [preciseOriginTS,
129         receivedCorrectionField,
130         syncRecTS]);
131     clockNotified := true;
132 );
133
134 public makeTimeAvailable : () ==> ()
135 makeTimeAvailable() ==
136     timeInfoAvailable := true;
137
138 public getTimeInfo : () ==> seq of real
139 getTimeInfo() ==
140     return [preciseOriginTS,
141             receivedCorrectionField,
142             cummulativeRateRatio];
143
144 public timeInfoRead : () ==> ()
145 timeInfoRead() ==
146     timeInfoAvailable := false;
147
148 -- Sets the slave port that corresponds to the current device
149 public setSlavePort: Slave ==> ()
150 setSlavePort(s) ==
151     slavePort := s;
152
153 -- Sets device ID

```

Appendix C. AVB case study VDM models

```

154 public setDeviceID: nat ==> ()
155 setDeviceID(id) ==
156     deviceID := id
157 pre deviceID = 0
158 post deviceID <> 0;
159
160 -- Gets device ID
161 public getDeviceID: () ==> nat
162 getDeviceID() ==
163     return deviceID
164 pre deviceID <> 0;
165
166 -- Receives the POTS
167 public receivePOTS: () ==> ()
168 receivePOTS() ==
169     preciseOriginTS := slavePort.receivePOTS();
170
171 public setLocalFreq : nat ==> ()
172 setLocalFreq(f) ==
173     freq := f;
174
175 -- Sets the frequency of the master device.
176 public setMasterFreq : nat ==> ()
177 setMasterFreq (f) ==
178     masterFreq := f
179 pre f <> 0;
180
181 -- Get the frequency of the master device attached to
182 -- the slave port.
183 public getFreq : () ==> ()
184 getFreq() ==
185     if (not isGM) then setMasterFreq(slavePort.getFreq());
186
187
188 -- Receives the follow up message sent by the master.
189 -- Reception is performed in one operation.
190 public receiveFollowUpMessage: () ==> ()
191 receiveFollowUpMessage() ==
192 ( let s : seq of real = slavePort.receiveFollowUpMessage() in
193   (
194     atomic (
195       receivedCorrectionField := s(1);
196       preciseOriginTS := s(2);
197       cummulativeRateRatio := s(3); );
198     if s(1) <> 0 and s(2) <> 0 and s(3) <> 0 then
199       followUpMessageReceived := true;
200   );
201 );
202

```

Device

```
203 -- Receives the cummulative rate ratio at the slave
204 -- port. Cummulative rate ratio is sent by the master
205 -- device attached to the slave port.
206 public receiveCRR: () ==> ()
207 receiveCRR() ==
208     cummulativeRateRatio:= slavePort.receiveCRR();
209
210 -- Calculates the neighbour rate ratio. This calculation
211 -- is performed at every slave device.
212 public calculateNeighbourRR : () ==> ()
213 calculateNeighbourRR () ==
214     if (not isGM) then
215         neighbourRateRatio := freq / masterFreq
216     else
217         neighbourRateRatio := 1
218 pre not isGM => (freq <> 0 and masterFreq <> 0)
219 post neighbourRateRatio <> 0;
220
221 -- Calculates the new cummulative rate ratio to be sent
222 -- to the attached slave unit.
223 public calculateCRateRatio : () ==> ()
224 calculateCRateRatio () ==
225     newCummulativeRateRatio := cummulativeRateRatio *
226                               neighbourRateRatio
227 pre neighbourRateRatio <> 0 and cummulativeRateRatio <> 0;
228
229 -- Calculates the new correction field to be sent to the
230 -- attached slave unit.
231 public calculateNewCorrectionField : () ==> ()
232 calculateNewCorrectionField () ==
233     if isGM then ()
234     else (
235         newCorrectionField := linkDelay + residenceTime *
236         newCummulativeRateRatio + receivedCorrectionField););
237
238 -- Receives a correction field sent by master unit.
239 public receiveCorrectionField : () ==> ()
240 receiveCorrectionField() ==
241     receivedCorrectionField := slavePort.receiveCorrectionField();
242
243 -- Sets the current device as a final device.
244 -- A final device will not have slave devices and
245 -- will not be acting as a master device.
246 public setFinalDevice: () ==> ()
247 setFinalDevice() ==
248     finalDevice := true
249 pre not isGM
250 post finalDevice;
251
```

Appendix C. AVB case study VDM models

```
252 -- Receives a synchronization sync event and
253 -- time stamps it.
254 public receiveSyncEvent: () ==> ()
255 receiveSyncEvent () ==
256 (
257   receivedSyncEvent := slavePort.receiveSyncEvent();
258   receivedSyncEvent := true;
259   updateTimeRef();
260 );
261
262 -- Updates the time reference when the time stamp operation
263 -- is invoked.
264 public updateTimeRef: () ==> ()
265 updateTimeRef() ==
266 (
267   --IO`print("\n Sync event --> timestamp generated");
268   syncRecTS := slavePort.getIngressTime();
269 );
270
271 -- Shows the device state
272 public showDevice : () ==> ()
273 showDevice () ==
274 (
275   IO`print("\nInfo for device: ");
276   IO`print(deviceID);
277   IO`print("\nIs GM:");
278   IO`print(isGM);
279   IO`print("\nIs final device: ");
280   IO`print(finalDevice);
281   IO`print("\nIs sync event received: ");
282   IO`print(receivedSyncEvent);
283   IO`print("\nSync event received at: ");
284   IO`print(syncRecTS);
285   IO`print("\nFrequency: ");
286   IO`print(freq);
287   IO`print("\nMaster frequency: ");
288   IO`print(masterFreq);
289   IO`print("\nLinkDelay: ");
290   IO`print(linkDelay);
291
292   IO`print("\nNeighbour Rate Ratio: ");
293   IO`print(neighbourRateRatio);
294   IO`print("\nRecCummulative Rate Ratio: ");
295   IO`print(cummulativeRateRatio);
296
297   IO`print("\nReceived CorrectionField: ");
298   IO`print(receivedCorrectionField);
299   IO`print("\nPrecise Origin TimeStamp: ");
300   IO`print(preciseOriginTS);
```


Environment

```
301 );
302
303 thread
304   while true do
305     (
306       timeStep();
307     );
308
309 sync
310   mutex(updateTimeRef);
311   mutex(timeStep);
312   mutex(notifyClock, receiveFollowUpMessage);
313
314   mutex(timeStep, receivePOTS);
315   mutex(timeStep, receiveCRR);
316
317   -- mutex(timeStep, receiveSyncEvent);
318
319 end Device
```

C.4.6 Environment

```
1 class Environment
2
3 types
4
5   public timeStimulus = seq of real;
6
7 instance variables
8
9   simSteps : nat := 3;
10
11   timeStimuli : seq of timeStimulus := [];
12
13   --public static timerRef : TimeStamp := new TimeStamp(5);
14
15 operations
16
17 public loadStimuli: () ==> ()
18 loadStimuli() ==
19 (
20   --loadStimulus([20, 30, 50]);
21   let i : seq of timeStimulus = [[40, 20, 10], [30, 80, 90]] in
22     for all e in set {1, ..., len i} do
23       loadStimulus(i(e));
24 );
```

Appendix C. AVB case study VDM models

```

25
26 public loadStimulus: timeStimulus ==> ()
27 loadStimulus(s) ==
28   timeStimuli := timeStimuli ^ [s];
29
30 -- Model simulation entry point
31 public run: () ==> ()
32 run() ==
33 (
34   IO`print("\nEndpoint device receiving time corrections");
35
36   loadStimuli();
37   IO`print(timeStimuli);
38
39   -- Simulation starts
40
41   start (Deployer`simClk);
42   start (Deployer`slave);
43   start (Deployer`clkSlave);
44   start (Deployer`siteSync);
45
46
47
48   duration(0) feedPhyMessage();
49
50
51   while Deployer`app.infoProvided() = false do
52     (
53       wait();
54     );
55
56   -- Shows the simulation results
57   Deployer`siteSync.showDevice();
58   Deployer`app.showReceivedTime();
59
60   IO`print("\nModel over\n");
61   IO`print("##### Unused time stimuli: ");
62   IO`print(timeStimuli);
63
64 );
65
66 public feedPhyMessage: () ==> ()
67 feedPhyMessage() ==
68 (
69   Deployer`channel.signalSyncEvent();
70   Deployer`channel.setFollowUp(hd timeStimuli);
71   timeStimuli := t1 timeStimuli;
72 );
73

```

PhyMessage

```
74
75 wait: () ==> ()
76 wait() ==
77     skip;
78
79 end Environment
```

C.4.7 PhyMessage

```
1  class PhyMessage
2
3  instance variables
4
5      -- sync event flag. Holds true value if a sync event
6      -- is put on the bus.
7      private syncEvent : bool := false;
8
9      -- Contents of the follow up message
10     private preciseOriginTS : real := 0.0;
11     private cummulativeRateRatio : real := 0.0;
12     private correctionField : real := 0.0;
13
14     -- Frequency sent by a master
15     private masterFreq : real := 0;
16
17     private fupSet : bool := false;
18
19 operations
20
21 -- Gets the contents of the follow up message.
22 public getFollowUp: () ==> seq of real
23 getFollowUp() ==
24     if not fupSet then
25         return []
26     else
27         return [correctionField,
28                 preciseOriginTS,
29                 cummulativeRateRatio];
30
31 -- Sets the contents of the follow up message.
32 public setFollowUp: seq of real ==> ()
33 setFollowUp(s) ==
34 (
35     setCRR(s(1));
36     setCorrectionField(s(2));
37     setPOTS(s(3));
```

Appendix C. AVB case study VDM models

```

38     fupSet := true;
39 );
40
41 -- Signals a Synchronization event
42 public signalSyncEvent: () ==> ()
43 signalSyncEvent() ==
44 (
45     --IO`print("\nMaster port --> Sync event sent");
46     syncEvent := true;
47     IO`print("Message in hardware buffer at: "); IO`print(time);
48 );
49
50 -- Check if there is a Synchronization event
51 public getSyncEvent: () ==> bool
52 getSyncEvent() ==
53 (
54     --IO`print("\nSlave port --> Sync event received");
55     return syncEvent;
56 );
57
58 -- #####
59 -- Getters and setters for the different fields in the
60 -- Physical Message.
61
62 public setMasterFreq : real ==> ()
63 setMasterFreq(f) ==
64     masterFreq := f;
65
66 public getMasterFreq: () ==> real
67 getMasterFreq() ==
68     return masterFreq;
69
70 public setCRR: real ==> ()
71 setCRR(crr) ==
72     cumulativeRateRatio := crr;
73
74 public setCorrectionField: real ==> ()
75 setCorrectionField(crf) ==
76     correctionField := crf;
77
78 public setPOTS: real ==> ()
79 setPOTS(pots) ==
80     preciseOriginTS := pots;
81
82 public getCField: () ==> real
83 getCField() ==
84     return correctionField;
85
86 public getPOTS: () ==> real

```

Port

```
87 | getPOTS() ==
88 |   return preciseOriginTS;
89 |
90 | public getCRR: () ==> real
91 | getCRR() ==
92 |   return cummulativeRateRatio;
93 |
94 | sync
95 |
96 |   mutex(getFollowUp, setFollowUp);
97 |   mutex(signalSyncEvent, getSyncEvent);
98 |
99 | end PhyMessage
```

C.4.8 Port

```
1 | class Port
2 |
3 | instance variables
4 |
5 |   -- Physical channel associated to the port
6 |   protected channel : PhyMessage;
7 |
8 |   -- Associated ID to the port
9 |   protected portID : nat;
10 |
11 |   protected clk : Clock;
12 |
13 | operations
14 |
15 | public timeStamp : () ==> ()
16 | timeStamp() ==
17 |   is not yet specified;
18 |
19 | public timeStep : () ==> ()
20 | timeStep() ==
21 |   is not yet specified;
22 |
23 |   -- Change the associated channel to the slave
24 | public changePhyChannel : PhyMessage ==> ()
25 | changePhyChannel(ch) ==
26 |   channel := ch;
27 |
28 | public getPhyChannel : () ==> PhyMessage
29 | getPhyChannel() ==
30 |   return channel;
```

```

31
32 -- Changes the associated ID to the port
33 public changePortID: nat ==> ()
34 changePortID(id) ==
35   portID := id;
36
37 end Port

```

C.4.9 Slave

```

1 class Slave is subclass of Port
2
3 instance variables
4
5   ingresTime : real := 0;
6
7   -- sync event flag. Holds true value if a sync event
8   -- is put on the bus.
9   private syncEvent : bool := false;
10
11   -- Contents of the follow up message
12   private preciseOriginTS : real := 0.0;
13   private cummulativeRateRatio : real := 0.0;
14   private correctionField : real := 0.0;
15
16   private followUpReady : bool := false;
17
18   public clk : Clock;
19
20   partialTime : real := 0;
21
22
23 operations
24
25 public checkSync : () ==> ()
26 checkSync() ==
27 (
28   duration(0) partialTime := time;
29   duration(0) syncEvent := channel.getSyncEvent();
30   if syncEvent = true then
31     duration(0) timeStamp(); -- Time stamp
32   );
33
34 public checkFollowUp: () ==> ()
35 checkFollowUp() ==
36 (

```

Slave

```

37  let i : seq of real = channel.getFollowUp() in
38  (
39    if len i = 0 then
40      followUpReady := false
41    else
42      (
43        atomic (
44          correctionField := i(1);
45          preciseOriginTS := i(2);
46          cummulativeRateRatio := i(3));
47      )
48    );
49 );
50
51 public timeStamp : () ==> ()
52 timeStamp() ==
53 ( duration(0) ingresTime := clk.getTime();
54   duration(0) IO`print("Received at time: ");
55   IO`print(partialTime);
56 );
57
58 public getIngressTime : () ==> real
59 getIngressTime () ==
60   return ingresTime;
61
62 -- Class constructor
63 public Slave: PhyMessage ==> Slave
64 Slave(phy) ==
65   channel := phy;
66
67 -- Receives the cummulative rate ratio
68 public receiveCRR: () ==> real
69 receiveCRR() ==
70   channel.getCRR();
71
72 -- Receives the correction field
73 public receiveCorrectionField: () ==> real
74 receiveCorrectionField() ==
75   channel.getCField();
76
77 -- Receives the precise origin time stamp
78 public receivePOTS : () ==> real
79 receivePOTS() ==
80   channel.getPOTS();
81
82 -- Receives the follow up message
83 -- Receives Precise Origing Time Stamp, cummulative rate ratio
84 -- and correction field in one invocation
85 public receiveFollowUpMessage: () ==> seq of real

```

Appendix C. AVB case study VDM models

```
86 receiveFollowUpMessage() ==
87   return [correctionField, preciseOriginTS, cummulativeRateRatio];
88
89 -- Checks if there is a synchronization event
90 -- available
91 public receiveSyncEvent: () ==> bool
92 receiveSyncEvent() ==
93   return syncEvent;
94
95 -- Gets the frequency associated to the master
96 public getFreq: () ==> real
97 getFreq() ==
98   return channel.getMasterFreq();
99
100 public timeStep: () ==> ()
101 timeStep() ==
102 (
103   if not syncEvent then
104     checkSync()
105   else
106     if not followUpReady then
107       checkFollowUp();
108 );
109
110 thread
111
112   while true do
113     (
114       duration(0) timeStep();
115       -- -- duration(1000) ()
116     );
117
118     --periodic(2E5, 0, 0, 0) (timeStep);
119
120 end Slave
```

C.4.10 TimeStamp

```
1
2 class TimeStamp
3
4 values
5
6 public stepLength : nat = 1;
7
8 instance variables
```


TimeStamp

```

9
10 currentTime : nat := 0;
11 wakeUpMap    : map nat to [nat] := {|->};
12 barrierCount : nat1;
13
14 operations
15
16 public TimeStamp : nat1 ==> TimeStamp
17 TimeStamp(count) ==
18   barrierCount := count;
19
20 public WaitRelative : nat ==> ()
21 WaitRelative(val) ==
22   WaitAbsolute(currentTime + val);
23
24 public WaitAbsolute : nat ==> ()
25 WaitAbsolute(val) == (
26   AddToWakeUpMap(threadid, val);
27   -- Last to enter the barrier notifies the rest.
28   BarrierReached();
29   -- Wait till time is up
30   Awake();
31 );
32
33 BarrierReached : () ==> ()
34 BarrierReached() ==
35 (
36   while (card dom wakeUpMap = barrierCount) do
37     (
38       currentTime := currentTime + stepLength;
39       let threadSet : set of nat = {th | th in set domwakeUpMap
40         & wakeUpMap(th) <> nil and
41         wakeUpMap(th) <= currentTime }
42       in
43         for all t in set threadSet
44         do
45           wakeUpMap := {t} <-: wakeUpMap;
46     );
47 )
48 post forall x in set rng wakeUpMap & x = nil or x >= currentTime;
49
50 AddToWakeUpMap : nat * [nat] ==> ()
51 AddToWakeUpMap(tId, val) ==
52   wakeUpMap := wakeUpMap ++ { tId |-> val };
53
54 public NotifyThread : nat ==> ()
55 NotifyThread(tId) ==
56   wakeUpMap := {tId} <-: wakeUpMap;
57

```

Appendix C. AVB case study VDM models

```
58 public GetTime : () ==> nat
59 GetTime() ==
60     return currentTime;
61
62 Awake: () ==> ()
63 Awake() == skip;
64
65 public ThreadDone : () ==> ()
66 ThreadDone() ==
67     AddToWakeUpMap(threadid, nil);
68
69 sync
70     per Awake => threadid not in set dom wakeUpMap;
71
72     mutex(AddToWakeUpMap);
73     mutex(NotifyThread);
74     mutex(BarrierReached);
75
76     mutex(AddToWakeUpMap, NotifyThread);
77     mutex(AddToWakeUpMap, BarrierReached);
78     mutex(NotifyThread, BarrierReached);
79
80     mutex(AddToWakeUpMap, NotifyThread, BarrierReached);
81
82 end TimeStamp
```

Representations and models used at different levels of abstraction

An adequate level of description should be provided at each abstraction level in order to be able to describe and analyse the relevant aspects of it. The aspects of interest at different abstraction levels vary considerably, therefore several modelling languages and methods, are required to create expressive and adequate representations at each level. Several authors [Wolf03, Waddington&06, Edwards&97] remark the fact that heterogeneous modelling is the most promising way of working with Hardware/Software systems. By using different modelling techniques in a systematic manner, clearly defining and understanding borderlines and interfaces between the representations, the best out of each modelling technology can be obtained. Besides defining the links between the models, it is responsibility of the modeller to apply the correct language at each level of abstraction, considering the kind and amount of details the model has to incorporate. For example, incorporating specific hardware details at an algorithmic level would not be relevant to discuss the system behaviour from a software perspective, further more, such a level of detail would be an inconvenient while understanding the software components of the system. On the other hand, specific hardware details are essential in a CAM model.

At an algorithmic level, languages like Unified Modelling Language (UML), Matlab or Specification Description Languages (SDL) are efficient in terms of expressiveness and precision. SystemC is performing better than the previous languages when it comes to model hardware implementations, furthermore, it is flexible and can be used across the interface level, the transaction level and the cycle accurate level. However, since it is an extension of the programming language C++, models tend to be mixed with implementation details.

In general, the lower the level of abstraction, the more details will be present. This implies more precision and complexity in the system representation. Taking as an example a hardware block, at the algorithmic level it might be represented by a black box that provides an interface to the application logic. This abstraction is isolating memory addresses from business logic through functions to read/write data to the interface. At an interface level it is relevant to know where exactly the data is placed and which procedure has to be followed in order to read/write data to the memory block. Details that were not relevant at an algorithmic level now are becoming important e.g.: concrete memory addresses, reserved areas... Figure D.1 shows part of this information and an adequate representation of it. Note that at this level time details are not relevant, details that will be incorporated in a TLM model when the current representation is refined.

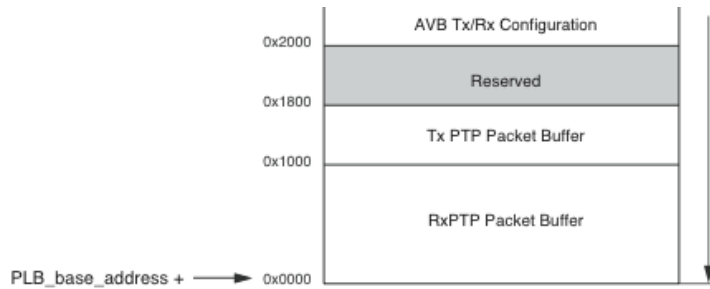


Figure D.1: Part of the memory map of an Audio Video Bridging (AVB) core interface. This memory map serves as an interface between the software and the hardware layers.

The description provided at the interface level is sufficient and precise enough for a programmer that is using the hardware, and is adequate for a hardware engineer to understand the interfaces that should be provided by the hardware layer. However a hardware engineer needs to represent more details while modelling a hardware block. In this case, what happens between clock cycles becomes relevant, and the information representation goes down to the bit level. At this level, specific Hardware Description Languages (HDL) like VHDL or Verilog are used. As an example of the level of detail that can be modelled with this languages, a VHDL representation of a tri-state buffer¹ is presented in D. This representation is useful for a hardware engineer aiming to implement the actual RTL layer. The architecture and behaviour are described at the bit level, e.g: on line 3 three bit variables are declared. The values these bits are holding during tristate operation are further specified in lines 10, 13 and 14.

```

1 entity test_tristate is
2   generic ( width : integer := 17 );
3   port ( en1, en2, en3      : in  std_ulogic;
4         inp1, inp2, inp3    : in  std_logic_vector (width downto 0);
5         tribus              : out std_logic_vector (width downto 0));
6 end test_tristate;
7
8 architecture rtl of test_tristate is
9 begin
10  tribus <= inp1 when (en1 = '1') else
11    (others => 'Z');
12  tribus <= inp2 when (en2 = '1') else
13    (others => 'Z');
14  tribus <= inp3 when (en3 = '1') else
15    (others => 'Z');
16  tribus <= ( others => 'H' );
17 end rtl;

```

It is possible to go even further and analyse the actual silicon implementation in the FPGA. In figure figure D.2, the used silicon area is shown. This representation of the system is critical for fine-tuning, floor planning and optimization. As it has been shown, multiple system views are offered by multiple modelling and representation technologies. Different representation techniques

¹Logic device able to output a high or low logical level or present a high impedance (disconnected) state

(in the end, modelling techniques) at the implementation level are widely accepted in industry. All of the above presented models, are providing relevant details at different levels of abstraction. Note that so far the models have been performed considering a single system. Besides this scenario, many problems are solved by systems that are communicating with other active systems, e.g. a network bridge, which implements certain business logic and communicates with other devices. This introduces an additional abstraction level, which is the Systems of Systems view. Certain communication and interaction requirements from the System of System point of view might be worth to consider while performing the Hardware/Software partitioning. Different partitioning schemes might have an impact from the distributed-real time perspective that should be considered in the development process.

Additional details on the properties the modelling languages should present will be provided in section 2.4. Further discussion on modelling technologies in the Hardware/Software co-design area will be exposed in chapter 4.




Figure D.2: Used silicon area in a programmed FPGA. Layout information can be obtained through specific tools provided by the manufacturer. Source: xilinx.com

MATLAB Based Methodology for Hardware/Software Validation

The official MATLAB based methodology is considering primarily three scenarios. In the first one, shown in figure E.1 MATLAB is used as a test bench providing on-the-fly stimulus to an HDL block, simulated in a specialized EDA tool and to Algorithms running in MATLAB. Later on, results can be analysed in the MATLAB tool.

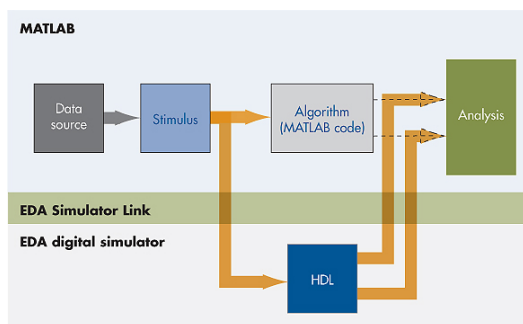


Figure E.1: MATLAB Simulink used as a testing infrastructure. Source: [Simulink]

In the second scenario, figure E.2, a MATLAB block is substituting an Hardware Definition Language block. The advantage of such an approach is that different functionality can be evaluated rapidly by implementing it in a high level programming language, without having to implement it in HDL. It must be remarked that this second simulation is taking place in the EDA simulator tool, and using MATLAB only as a functionality provider.

The last scenario, shown in figure E.3, is a combination of the previous two. Simulink is used to provide stimulus and to implement some of the HDL blocks. Part of the functionality is run in the EDA simulator and the rest in Simulink blocks. Finally, simulation results are analysed in Simulink.

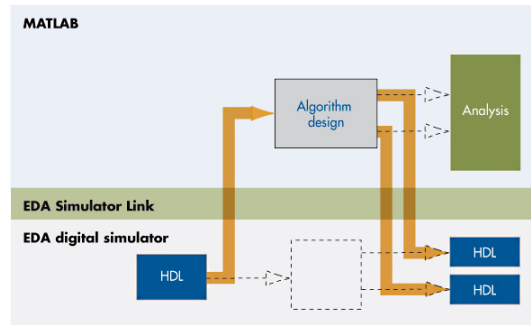


Figure E.2: MATLAB Simulink combining an algorithm block and an HDL implementation. Source: [Simulink]

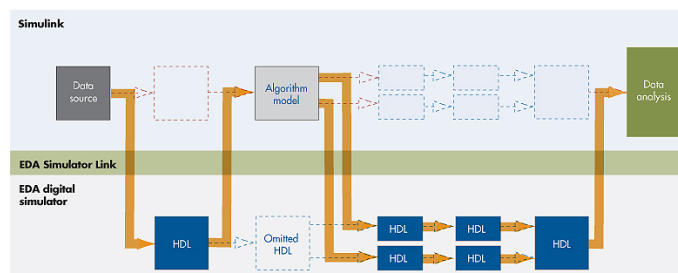


Figure E.3: Combined test approach for HDL and Simulink simulations. Source: [Simulink]