

Metropolis Design Guidelines

Alessandro Pinto
University of California at Berkeley
545P Cory Hall, Berkeley, CA 94720
apinto@eecs.berkeley.edu

September 14, 2004



Copyright © 2001-2004 The Regents of the University of California.
All rights reserved.
UCB/ERL M04/40

Contents

Contents	2
1 Introduction	5
1.1 Platform-based design of Multimedia Systems	5
1.2 Metropolis design environment	8
1.3 Audience and Organization	10
1.4 Acknowledgments	11
2 Basic notions on modeling	13
2.1 Processes, Interfaces and Media	13
2.2 Netlists	18
2.3 Quantity managers	19
2.4 Constraints	20
3 Developing platforms for functional description	23
3.1 Models of computations for describing functions	24
3.2 Function platforms use-case	25
3.3 Architecture of a model of computation	27
3.4 YAPI: Y-Chart Application Programming Interface	29
3.5 Muti-rate Synchronous model of computation	34
4 Developing platforms for architectural description	37
4.1 Architecture platforms use-case	38
4.2 Architecture platforms development	40
4.3 The single-processor-single-memory architecture	45

5	Mapping a function onto an architecture	51
5.1	Mapping functions onto architectures	51
5.2	Describing mappings using metamodel synchronization constraints	53
A	Platform-Based Design example	55
B	Tagged Signal Model Definitions	59
	Bibliography	61

Chapter One

Introduction

The Metropolis design environment provides an infrastructure for designing embedded systems. Metropolis is particularly suited to model *heterogeneous* systems at *different levels of abstractions* and it is constantly under development with the idea in mind of supporting the platform-based design [] principle. The main objective, in fact, is not merely to provide a simulation/verification/synthesis environment for design electronic systems at a specific level of abstraction, but rather to develop and infrastructure that is flexible enough to be used for developing design flows for different application domains.

This ambitious goal is pursued by supporting design principles that are independent on the specific application. Metropolis follows two basic principles: *orthogonalization of concerns* and *platform-based design* [7]. The latter can be applied in many fields of engineering. In this design guide, we use a multimedia design flow as a representative example.

1.1 Platform-based design of Multimedia Systems

A simple example of platform-based design applied to the logic synthesis flow is described in appendix A. In this chapter we give an overview of the method for multimedia applications.

1. INTRODUCTION

Going from the denotational description of a function down to its implementation is a very hard problem. The number of possible design choices makes the design space too large to be explored efficiently. The big design gap can be subdivided into smaller steps by introducing a stack of platforms, each dedicated to the exploration of the design space along few directions (figure 1.1).




Figure 1.1: Platform stack for multimedia designs

When using a platform-based design methodology the first important step is to define those level of abstractions. Consider, for instance, the case of multimedia systems design as shown in figure 1.1. Starting from the denotational description of the function F the first important decision is to select a suitable platform (that at this level of abstraction is usually referred to as a model of computation) to describe F . The only property that we are interested in at this level is correctness. Kahn process networks (KPN) [6] is a convenient model of computation usually adopted as the first platform for modeling multimedia systems. The platform components are processes running on their own threads and

communicating through unbounded FIFOs with blocking read and non-blocking write semantics.

Mapping F onto the Kahn process networks platform implies re-expressing the denotational algorithm as interconnection of concurrent processes and unbounded FIFOs. For instance, an FIR filter can be described by interconnecting adders and multipliers together. At this level of abstraction we can check if our algorithm is correct without caring about deadlocks due to resource limitations like FIFO boundedness. Also, since processes are totally concurrent and write is non-blocking, processes don't have to compete for shared resources. The designer can then explore the maximum amount of concurrency (parallelism) in the functional description without being constrained by resource limitations. However, the function is too abstract to be implemented on a real architecture which has resource limitations, e.g. memory size.

The next level of abstraction is represented by the TTL [4] (task transaction level) platform. It is still composed of processes running on their own threads, but the communication among them is implemented with bounded FIFOs. Mapping a KPN function onto a TTL platform is a simple one to one mapping that can be done by a direct refinement of each unbounded FIFO into a more complicated channel as we will see in the next section.

At the TTL level of abstraction, we are concerned with memory size minimization. Each communication channel is parameterized by token-size and maximum number of tokens that it can contain. Depending on the interleaving policy between writer and reader of the same channel, buffer size can be reduced at the expense of a more frequent context switching between the two processes. The memory usage/switching overhead trade-off is explored at this level of abstraction.

Finally the TTL description is mapped onto a real micro-architecture, for instance a single processor architecture. At this level of abstraction we are concerned with memory allocation and organization and task scheduling. Task scheduling is needed because at this level we have a computational resources limitation (only one processor in this example). A real time operating system works as an adaptation layer between the high level of concurrency of the TTL platform and the sequential operations of a single processor architecture.

The result of this mapping is the implementation of the original function on a target micro-architecture that satisfies all the original constraints.

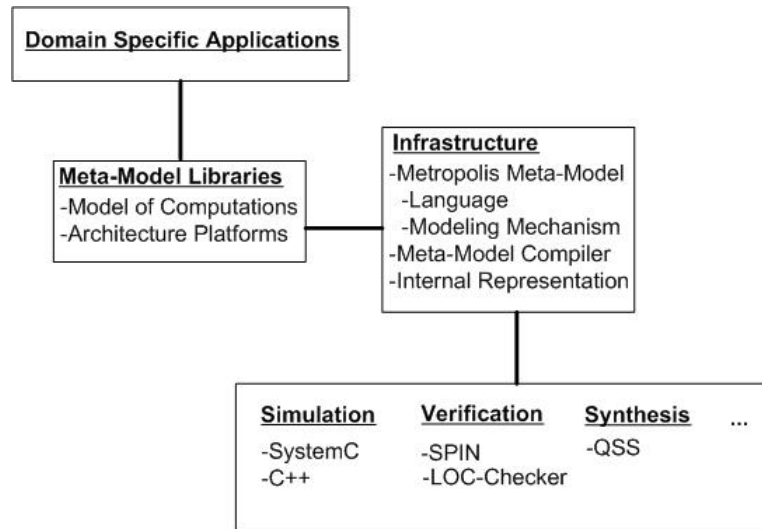


Figure 1.2: The Metropolis framework

1.2 Metropolis design environment

The idea behind the Metropolis framework is to have an unambiguous common representation of designs and a set of tools that can interpret the common representation, manipulate it and generate a modified description using the same representation [3].

The Metropolis framework is pictorially represented as in figure 1.2

The Metropolis infrastructure is the core of the framework. The infrastructure is composed of a language for the design description that is called the Metropolis Meta-Model (MMM). An MMM program is compiled into an internal representation that retains all the semantic and syntactic information of the original MMM program.

A given application could be directly described using the MMM. However, a set of platforms are provided to designers to facilitate the system description. There are two kinds of platforms: model of computations and architectures. As described in chapter 3, a model of computation is presented as a set of basic classes that the user extends to customize the behavior of processes and obtain the functional description. As described in chapter 4, an architecture is presented as a set of services and their legal compositions.

Given an application domain, the user selects a suitable model of computation from the available set of libraries. For instance, Kahn process networks (KPN) are widely used in modeling multimedia applications. The KPN library in Metropolis provides processes and media. A process is a basic class with an empty behavior. The user extends this class by adding ports for external communication and overriding his thread to describe the set of behaviors that belonging to the process. Media are unbounded FIFOs that the user has only to instantiate and connect to processes in order for them to communicate. The description of the entire system is still done in MMM but the library considerably simplifies its description.

The user now selects an architecture which could be available already or has to be assembled from a set of components. The architecture is also described using the MMM language.

Finally a mapping is obtained by enforcing a synchronization of function and architecture. Each action in the function side is correlated with a function on the architecture side using synchronization constraints. Semantically this means that actions in the architecture have to follow the execution of the same action in the function side and its a way of taking the intersection of the two in the common semantic domain.

Function, architecture and mapping are all described using the MMM. Since all back-end tools are developed upon the internal data representation, they can be applied to the function, architecture and mapping netlist. It is possible for instance to simulate the architecture without any function mapped onto it.

Given a metropolis netlist, it can be parsed and compiled to generate the internal representation using the command *metacomp*. Compilation is not the more interesting step but elaboration is probably the most useful. Elaboration can be invoked by *metacomp -elaborator ;topnet; jfiles.mm;* where *topnet* is the top netlist. This phase compiles the source files and also run the constructor of each component to determine its initial state. It also checks interfaces, types and modifiers.

A number of other tools are available in the distribution. The most used is certainly the SystemC simulator that after elaboration generates an equivalent description of the original specification in the SystemC language. It also generate a makefile to compile the C++ program and generate an executable. The SystemC back-end can be run using the command line *metacomp -systemc -top ;topnet; jfiles.mm;*

Also, a shell *metroshell* is provided to compile and elaborate an MMM program. This shell provides APIs to browse the design and apply back-end tools to it.

1.3 Audience and Organization

The design guideline is intended to give directions to users that want to design embedded systems in the Metropolis framework. This document touches two aspects: what the user requirements are and how developers should satisfy them. This document does not describe syntax and semantics of the MMM, backend tools, or any other implementation related issues. We target this guidelines to two class of users: system developers that have to describe applications, select architectures and perform mapping, and libraries provider that have to provide the necessary infrastructure to make the system developers job easier. This guidelines is then an overview of *how* the Metropolis framework, and specifically the MMM, should be used.

The Metropolis guidelines document is organized as follows:

- **Chapter two** describes the basic components used to describe a system. Processes, media, quantity managers and netlists are introduced through a simple example and a brief description of their usage is given.
- **Chapter three** shows how a model of computation library should be used and consequently how it should be developed. Two examples are given: KPN and a multirate synchronous domain that are both provided by the Metropolis release.
- **Chapter four** shows how an architecture platform is assembled and used. It describes how the services offered by an architecture should be exposed to a user and consequently how all components should be developed and described. An example of simple architecture is given.
- **Chapter five** described the mapping phase. This chapter only describes the strategy that we use in the Metropolis framework to describe a mapping and also give a small example.

1.4 Acknowledgments

This work was supported in part by the following corporations: Cadence, General Motors, Intel, Semiconductor Research Corporation (SRC), Sony, STMicroelectronics; and the following research projects: NSF Award Number CCR-0225610 and the Center for Hybrid and Embedded Systems (CHESS, <http://chess.eecs.berkeley.edu>), The MARCO/DARPA GigaScale Silicon Research Center.

The Metropolis project would also like to acknowledge the research contributions by: The Project for Advanced Research of Architecture and Design of Electronic Systems (PARADES, <http://www.parades.rm.cnr.it/>) (in particular Alberto Ferrari), and Politecnico di Torino, Carnegie Mellon University, University of California, Los Angeles, University of California, Riverside, Politecnico di Milano, University of Rome, La Sapienza, University of L'Aquila, University of Ancona, Scuola di Sant'Anna, University of Pisa, .

Thanks to Trevor Meyerowitz, Luciano Lavagno and Felice Balarin for their useful feedback. Thanks to Alberto Sangiovanni-Vincentelli and Roberto Passerone for the theory of Platform-Based Design.

Chapter Two

Basic notions on modeling

The Metropolis Meta-Model (MMM) language provides basic components that are used to describe a system. These components represent computation, communication and synchronization which are three basic ingredients needed to define a model of concurrency. This chapter uses a simple producer consumer example to introduce the basic objects defined in the MMM language. For a detailed explanation of MMM syntax and semantics refer to [2].

2.1 Processes, Interfaces and Media

Consider a simple producer consumer example. The producer generates a sequence of integers in ascending order starting from 1 and communicates them to the consumer whose it just a sink.




Figure 2.1: A simple producer consumer example

Producer and consumer are two objects that execute an algorithm. More specifically, the producer follows a certain number of steps to implement the original specification. These steps are a set of “actions”. Objects of this type are defined as processes in the MMM language. A process runs on his own thread. A thread is a sequence of actions that can be thought of as instructions, sub-instructions (in which an instruction can be decomposed), function calls and awaits (for a formal definition of actions please refer to [2]).

The code implementing the producer follows:

```
process producer {
  port ReadInterface r;
  port WriteInterface w;
  parameter int N_writes;
  int v;
  producer(String n, int nw) {
    super(n);
    N_writes = nw;
  }
  public void thread(){
    int j;
    for (int i=0; i<N_writes; i++){
      v = r.read();
      j = i + 1;
      out.write(j);
    }
  }
}
```

The process has two ports and one parameter, which is specified at instantiation time. Note that ports don’t have direction, but have a type associated with them. The type of a port is an interface which declares services that can be called by the process.

Figure 2.1 also shows a possible connection of the process to the rest of the system. Processes cannot connect directly to other processes but the interconnection has to go through a medium which has to define (i.e. implement) the services declared by the interface associated with the ports that access the medium.

An interface is an “abstract” class which just declares a set of functions with their associated signatures. The *WriteInterface* for instance is specified as follows:

```
interface WriteInterface extends Port {
  eval void write(int data);
}
```

The interface declares a service *write* which takes an integer *data* as the only parameter. It does not define the function so it is impossible to say what is the semantics of this function. If a process has a port of type *WriteInterface* it has to implement the *write* service because the process is going to use it. The interface is then an agreement on what is offered and what is required in the sense that a process can only use services declared in the interface and on the other hand a medium connected to that port has to provide those services.

The syntax used in MMM is very close to Java. A process, like any other object in the MMM, has a constructor which in this case takes a parameter *nw* indicating the number of integers to be produced in ascending order starting from one.

A process specifies a *thread* function. In this example, producer reads a triggering signal from port *r* and writes an integer to port *w*.

From the outside, the process behavior is observed as a trace of read and write actions, or better yet as a trace of services calls to the media connected to its ports. More precisely the behavior is a sequence of events which are the begin and end of each action. Looking at the process thread, the behavior is more complex. After reading the trigger, an internal variable (belonging to the state of the process) is assigned to the result of a sum and written to the output port. The behavior, then, not only includes read and write but also begin and end events of the assignment and sum actions.

A process is then an object that generates a sequence of events. Each process in a system evolves by executing one event after the other. At each step (which is formally described in [2] in terms of a global execution index), each process in the system executes one event. This is the semantics of a synchronous language but a special event called *NOP* is defined in such a way that it can always be interleaved between to events of a process. The *NOP* event correspond to the stalling of a process making it possible to implement asynchrony. For instance, if producer wants exclusive access to medium *M2* we can force consumer to execute the *NOP* event while producer is writing a data.

The juxtaposition of events from all processes is an event vector. At each step, there is a set of event vectors that could be executed to make a transition from the current state to the next state. If there are no scheduling constraints specified then the choice among all possible transitions is done non-deterministically.

2. BASIC NOTIONS ON MODELING

A communication medium is an entity that implements services. It does not have a thread of execution but rather inherits the thread from the process that uses the services. A communication medium implements a protocol to exchange information between processes. Separation between processes and media follows the principle of orthogonalization of concerns described in [7].

A media implements an interface which is a declaration of a set of services and gives them a semantics. This example shows a simple channel that implements two services: read and write.

```
medium Channel implements ReadInterface, WriteInterface {
    int[] storage;
    int space, n; reading; writing;
    int length;

    public IntM(String name, int nelement) {
        n = 0;
        space = nelement;
        storage = new int[nelement];
        length = nelement;
        reading = 0;
        writing = 0;
    }

    public update void write(int token) {
        await {
            (space > 0; ; ) {
                space = space - 1;
                n = n + 1;
                storage[writing] = w;
                writing = writing + 1;
                if (writing == length) writing = 0;}}}

    public update int read() {
        int _retval = 0;
        await{
            (n > 0; ; ) {
                n = n - 1;
                space = space + 1;
                _retval = storage[reading];
                reading = reading + 1;
                if (reading == length) reading = 0; }
            return _retval;}}}
```

The await statement can define parallel critical sections. Each critical sec-

tion has a premises (*guard; testlist; setlist*). In this example we only have a guard condition that is a boolean expression. If the guard is true that the critical section can be entered. The meaning of *testlist* and *setlist* will be explained later through an example

The channel has a limited storage space. A reading process will be blocked until there is at least one token in the FIFO while a writing process will be blocked until there is at least one token space in the FIFO. Note that a channel implementing the same interfaces but using a different implementation could change the communication semantics. For instance, we could implement a non-blocking write operation by just overwriting the current data.

In this example, there is no restriction on the simultaneous access to the storage variable in the medium. A reader and a writer can access the medium at the same execution step and modify the internal state simultaneously. When this situation occurs, the final value of the variables is not known. In order to avoid data corruption, a synchronization method between the two processes has to be implemented. The synchronization protocol can be described directly in the medium (using the set lit and test lit in the await statements) or using an external scheduler that we call quantity manager.

We could use an empty interface to act as a semaphore and include it in the await test and set lists as follows:

```
interface Sem {};  
medium Channel implements ReadInterface, WriteInterface, Sem {  
  
    ...  
    public update void write(int token) {  
        await {  
            (space > 0;Sem;Sem) {  
                ...  
            }  
        }  
    }  
  
    public update int read() {  
        await{  
            (n > 0;Sem ;Sem) {  
                ...  
            }  
        }  
    }  
}}
```

We can imagine that each interface has a flag associated with it. Each time a critical section is entered, the flags of all the interfaces in the set list are raised. Before entering a critical section, though, not only the guard has to evaluate to true but all interfaces in the test list have to have their flags not raised. In our implementation, if the producer is executing a write operation it means that the *Sem* interface was flagged and hence a consumer will be forced to execute the *NOP* event. In this example the await statement is used to synchronize (schedule) the two process. The same result can be achieved by using another object called quantity manager as explained in section 2.3

2.2 Netlists

A netlist is an object used to instantiate and connect other components like processes, media and other netlists. A netlist has a constructor where all components are instantiated and connected. An API is provided to add a component to a netlist and to connect a port of an object to a medium. In our case for instance the netlist code looks as follows:

```
netlist ProducerConsumer {
    public ProducerConsumer(String n, int nw){
        super(n);
        producer p = new producer(``Prod``,nw);
        consumer c = new consumer(``Cons``);
        controller ctrl = new controller(``Contr``);
        Channel ch = new Channel(``CommCh``)
        addcomponent(p,this,``ProdInstance``);
        ...
        connect(p,w,ch)
        ...
    }
}
```

An object is created as in Java using the keyword *new*. Even if an object is in memory it has to be added to the netlist before using it. The function *addcomponent* adds the object to the netlist provided as second argument (the current netlist in this case). Finally it is possible to connect ports to medium using the function *connect*.

2.3 Quantity managers

Quantity managers are used to assign tags to events. A tag is an abstract quantity from a partially order set. Time, for instance, is a real number so in this case the set of tags is totally ordered. When an event has to be tagged with a quantity, an explicit request is made to the manager of that quantity. Due to concurrency of processes, multiple requests can be issued to a quantity manager that has to resolve them and schedule the processes in order to satisfy the ordering relation on the set of tags. Consider the simple producer/consumer example of figure 2.2. We want




Figure 2.2: A quantity manager

to make sure that internal variables are not accessed simultaneously by a reader and a writer. We can use a quantity manager to schedule the two processes:

```
quantity RWManager implements QuantityManager {
    port StateMediumSched reader;
    port StateMediumSched writer;

    public RWManager(String n){
        super(n);
        _pending = new ArrayList();
        //Constructor code }

    public eval void request(event e, RequestClass rc){
        _pending.add((Object)rc.clone());}

    public update void resolve(){
```

2. BASIC NOTIONS ON MODELING

```
        //For all pending requests
        // order them depending on the tags ordering
    }

    public update void postcond(){
        //Do the first event
        //Don't do the others
        _pending.clear();
    }

    public eval boolean stable(){
        return true;
    }

    ArrayList _pending;
}
```

The set of tags contains four elements: br, er, bw, ew . Each time the read service is called, before modifying the internal state it makes a request to annotate the begin of the read function with the tag br . Also, before exiting the read function, another request is made to annotate the end of the read function with the tag er . A similar set of requests is generated by the write service.

The quantity manager specifies an order on this set of tags that could be, for instance, $br < er, bw < ew, bw < br$. The last inequality says that in the case of simultaneous access of a reader and a writer, the read is executed first.

In general, quantity managers are used to adapt two different concurrency models. On one side, the set of processes are fully concurrent and can freely execute the sequence of events defined in their threads. On the other side, we have a specific process scheduling mechanism that we want to implement in order to give our program the semantics that we want. For instance, in a discrete time model, time is a totally ordered set of tags and hence a process that wants to execute an event at time t_i has to wait for other processes that are executing events at time $t_j < t_i$.

2.4 Constraints

An important part of the design specification is represented by constraints. They are declarative formulas that are used to specify property that the implementation have to satisfy. In our simple example, we might want

to declare that the time needed to complete a write operation has to be less than a certain quantity T .

Two kind of formulas are provided by the MMM language: linear temporal logic (LTL) and logic of constraints (LOC). The syntax is explained in [2].

Constraints are propagated down while marching towards the implementation and have to be constantly checked at each level of abstraction. A back-end tool called *LOC – Checker* is provided in the framework to verify that LOC formulas are satisfied.

At the functional level, where performance annotation are not available yet, constraints are uninterpreted formulas.

Chapter Three

Developing platforms for functional description

A design process always starts with the idea of what the system is supposed to do. In our terminology this is the *function*. It can be described, informally, using natural languages like English, or its description can be more formal. A denotational definition of the function usually involve formulas describing the outputs in terms of inputs. A maximum likelihood estimator, for instance, can be described as the solution of a maximization problem:

$$\max \ln P(r|c)$$

where r is the vector of the received bit stream and c is the vector representing the code. The description of the algorithm is denotational and only tells us what the system should do.

Practical and implementable algorithms exists to solve this problem. For example, the Viterbi decoding algorithm is an interconnection of blocks whose final result is an approximation of the maximum likelihood estimation. Being an approximation means that the result is different from the denotational description by a quantity e . If nothing has been stated about the error, then it could be any number but what we really want is that the implementation follows as close as possible the denotational description, where “as close as possible” is estimated by the error

e. Usually another important part of the functional description is a set of *constraints* that an implementation of the algorithm is subject to. In our example, for instance, a bound on the maximum admissible error can be set as constraint.

The denotational function is then mapped onto a platform that allows its description in a more “structured” way, usually as interconnection of sub-functions. Precise rules have to be imposed on this structure in order to have a non ambiguous description, or better, we need a uniquely defined way of interpreting the model that gives all and only the system’s execution traces.

The set of rules define the semantics of a *model of computation* (MoC). An MoC is defined in terms of how computation, communication and coordination must be carried out in a structured interconnection of objects.

This chapter gives few design guidelines for building an MoC (that we shall call functional platform in the sequel) in the Metropolis framework and introduces some examples of models of computation.

3.1 Models of computations for describing functions

This section is divided in three parts. This introduction explains briefly what are the properties characterizing a model of computation. The second shows how a user wants an MoC to be exposed to her/him in order to make the functional description easier. The third part is how a developer should build a platform and expose a model of computation to the users.

We will use the tagged signal model [8] (TSM) formalism (whose basic definitions are given in appendix B) as a denotational framework for stating properties about a model of computation.

Informally speaking, a model of computation can be defined by the set of values V , the set of tags T and the ordering relation on the tags \leq , the legal processes P and their communication semantics.

For instance, Kahan process networks are characterized by a set of tags that is partially ordered (it is an untimed model) and, since a com-

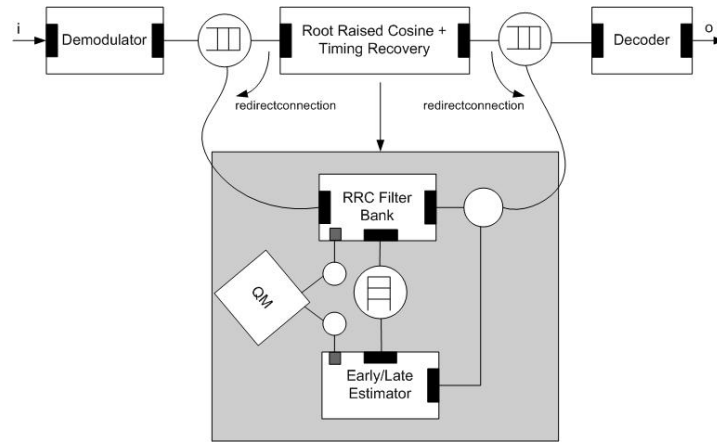


Figure 3.1: Block diagram of a digital receiver

munication channel has a first-input-first-output semantics, each signal is totally ordered.

An MMM process is used to describe all the possible behaviors of a process P while media are used to enforce a communication semantics. An order on the set of tags can be enforced by using a quantity manager which has the task of assigning tags to events. Consequently a quantity manager has to decide the scheduling of the process in order to satisfy the ordering of tags.

3.2 Function platforms use-case

Description of applications is considerably simplified when a suitable model of computation is chosen. Depending on the application domain, the first step in a design flow is to choose a natural function platform where the description and interconnection of functions is easy to carry out. For instance, digital signal processing algorithms are naturally expressed using a model where computation is done by blocks called actors that communicate through FIFOs.

Consider the following example. We would like to design a digital receiver which is the cascade of a digital demodulator, a root raised cosine filter, a timing recovery loop and a decoder. The block diagram is shown in the top part of figure 3.1.

3. DEVELOPING PLATFORMS FOR FUNCTIONAL DESCRIPTION

The system input signal *i* is a sequence of totally ordered events. As a first choice we select the Kahan process networks [6] (KPN) model of computation. We would like to describe the system as follows.

```
netlist DigitalReceiver {
  public DigitalReceiver(...) {
    //components instances
    StimuliGen s = new StimuliGen(...);
    Demodulator demod = new Demodulator(...);
    RRCTiming rrct = new RRCTiming(...);
    Decoder dec = new Decoder(...);
    //add components to this netlist
    addcomponent(s,this,"sinstance");
    ...
    //Interconnections
    FIFOinstance(s.out,demod.in);
    FIFOinstance(demod.out,rrct.in);
    FIFOinstance(rrct.out,dec.in);
  }
}
```

We first instantiate all components and add them to the current netlist and finally we interconnect them. Ideally we would like to have a way of specifying a connection by calling a function and passing the source and destination ports as parameters. In our example, this function is *FIFOinstance(port src, port dest)* that will instantiate a FIFO channel and connect the two ports to it. If the domain that we are using is KPN, there is no need for scheduling the processes since the blocking read and non-blocking write communication semantics will take care of the process synchronization. As result, the platform will not include a quantity manager or other synchronization constraints.

Consider now the refinement of the timing recovery loop shown in the bottom part of figure 3.1. The domain that we want to use in this case has different properties from the previous one. In particular we want to enforce a specific scheduling of the two processes. The sequence of reaction should be an infinite sequence of the RRC filter bank and the error detector. Using a quantity manager for ordering the events is the right solution. Every time a process want to read its inputs and compute its outputs it has to ask the quantity manager first. This request is basically asking to annotate the read-compute-write sequence of events.

The quantity manager, based on its pending queue of events, enforces the right scheduling deciding which process can proceed and which has to wait

A user of this model of computation would like to describe the new netlist as in the previous case without having any concern about the quantity manager and synchronization schemes.

A refinement of the original system is obtained by using the keyword *refine* which is explained in [2]. It is also necessary to redirect the original connection to the refining netlist. In this particular case, the communication semantics of the two domains is the same so a direct connection is possible.

However, a complex system could require the use of an heterogeneous model of computation because it spans multiple application domains. A communication system, for instance, is the interconnection of the radio sub-system to the base-band sub-system which is in turn connected to the data-link sub-system. Each of them has very different properties and a single model for their description is not the right solution. The radio sub-system is best described in a continuous time domain while the base-band sub-system is basically a dataflow kind of application. A direct connection of the two is not possible but a specific interface has to be designed to transform a dataflow signal into a continuous time one.

A user of the framework would like to have a set of interfaces between models that are available to be used. Moreover, each interface should be parameterizable so that the impact of translation of one signal into the other can be evaluated at this level of abstraction.

3.3 Architecture of a model of computation

A model of computation is provided as a library of components that a user extends and interconnects to describe a function. Section 3.2 introduces the user expectations which will determine the design guidelines for platform developers.

When developing a new platform we focus on the components to be exposed to the user, the way of customizing their behavior and a way of interconnecting them. The basic components to develop are shown in figure 3.2:




Figure 3.2: Components of a platform

- Processes for describing computation. A base class has to be provided together with a set of API to access ports values. After extending this class, a user has only to override a method to describe the process behavior.
- Media for describing communication between processes. The communication semantics is usually fixed in a model of computation, so a medium should not be touched by the platform user. However some parameters could be exposed to configure the medium.
- Quantity Managers for enforcing a scheduling policy of processes. Scheduling is also a property of a model of computation and should be out of the user control. The instantiation and interconnection of quantity managers should be totally transparent to the user. For instance, in a synchronous model of computation all events in an event vector have the same tag and even if this property is known to the MoC user, she/he should not be aware of how this property is enforced.

In the Metropolis framework, platforms for functional specification are provided as packages. All platforms reside in the repository tree under the directory *metro/lib/metamodel/plt*. It is good practice to create a new directory for each platform under development. For instance, a dataflow platform will reside under *metro/lib/metamodel/plt/dataflow*.

The rest of the chapter shows two examples of platforms available in the Metropolis framework. The first one does not use quantity managers while the second one, whose semantics has the notion of computation steps, needs a synchronization of the processes execution.

3.4 YAPI: Y-Chart Application Programming Interface

For a complete description of the YAPI semantics refer to [5]. Briefly, YAPI is a KPN model of computation where a non-deterministic choice has been added. More precisely, a process is allowed to check on the number of tokens on an input channel which is a way of implementing non determinism in a model of computation.

Processes communicate through unbounded FIFOs. Processes can read from and write to communication channels using two functions: *T[] read(int n)* and *void write(T[] data)* where *T* is the data type. It is possible to parameterize the functions with respect to the data types by using templates but another valid option is to define an interface that all valid data types have to implement. We will use the first option because it is more natural for designers to use templates that are widely used in C++.

The first thing to define are the interfaces:

```
package metamodel.plt.yapitemplate;
//interface to check fifos
public interface yapiinterface extends Port {
    eval boolean checkfifo(int n,int dir);}
//write interface
template(T)
public interface yapioutinterface extends yapiinterface {
    update void write(T data);
    update void write(T[] data,int n);
    update void write(T[][] data, int n, int m);}
//read interface
template(T)
public interface yapiininterface extends yapiinterface {
    update T read();
    update void read(T[] data,int n);
    update void read(T[][] data,int n,int m);}
```

We define a service to check how many tokens are present in the FIFO. Then we define a writing interface that declares three services for writing

3. DEVELOPING PLATFORMS FOR FUNCTIONAL DESCRIPTION

data: one for writing matrices, one for vectors and one for single data. Finally we define a reading interface for reading data. This is sufficient for a designer that don't really care about the services implementation but only about their signature and semantics. A communication channel is implemented as follows:

```
template(T)
public medium yapichannel implements yapiininterface-<T>- ,
                                     yapioutinterface-<T>- ,
                                     rdi,wri,cki {

    //Definition of internal buffers
    ...
    //Constructor
    public yapichannel(String n,int isize){
        super(n);
        ...}

    public update void write(T[] data,int n){
        await{
            (true;this.rdi,this.cki;this.wri){
                int i;
                if ((ntokens + n) >= size){ //resize internal buffer
                    ... }
                else { //OK we can write directly
                    for(i= 0; i< n ; i++){
                        FIFO[(wp+i)%size] = data[i].clone();
                        wp = (wp + n) % size;
                        ntokens = ntokens + n;}
                    }}}

    public update void read(T[] data,int n){
        await{
            (ntokens > n-1;this.wri,this.cki;this.rdi){
                for(int i=0 ; i<n ; i++){
                    data[i] = FIFO[rp];
                    rp = (rp +1)%size;
                }
                ntokens = ntokens - n;}
            }}}}
```

The channel is defined as a medium that implement all the interfaces shown before. It also implements some dummy interfaces that are used only to prevent data corruption. Simultaneous access to the same data is prevented using an await statement that before entering a critical section

checks if another process is using one of the dummy interfaces. Since the write is non-blocking, the guard condition on the corresponding await statement is always true. It might happen that the FIFO has to be re-sized because there is no more space to write. This method is a way of implementing an unbounded FIFO.

The read function instead, has a guard condition that checks the number of tokens in the FIFO before executing the read operation. If the number of tokens is less than requested, the calling process is blocked.

The YAPI platform process offers a function to select non-deterministically a port among a set of ports where a read or write operation can eventually be completed. This function is useful for handling inputs coming from a user. Reading from these inputs will block a process execution until a new token arrives. It is important to provide a way of continuing the execution if there are no tokens available on that channel. This situation occurs for instance when an input is used to change a process behavior during execution. The YAPI platform implements a process as follows:

```
template(T)
public process yapiprocess{

    public yapiprocess(String n){
        super(n);
    }

    public int select(ArrayList ports,int[] tokens){
        await{
            (atLeastOnePortEnabled(ports,tokens);){
                return selectOnePort(ports,tokens);
            }
        }
    }

    boolean atLeastOnePortEnabled(ArrayList ports,int[] tokens){
        //OR of all checkfifo
    }

    int selectOnePort(ArrayList ports,int[] tokens){
        //non-deterministic selection of one port
    }

    void thread() {
```

```

        execute();}

    public void execute(){}

}

```

The select function is provided in the process implementation. It takes an array of ports and an array of integers where for each port $p[i]$ we request to read/write $n[i]$ tokens. Select first checks whether there is a least one ports where read or write will not block the process execution. If this condition is true, it selects one among these ports non-deterministically.

The thread function calls an execute function that the platform user override to describe the process behavior.

As an example of usage of the YAPI platform, consider the cascade interconnection of a producer, a filter and a consumer. The filter block implements two algorithm to filter the producer data. A controller decides which filtering algorithm to use (figure 3.3).




Figure 3.3: Block diagram of the producer-filter-consumer example

The filter code is the following:

```

import metamodel.plt.yapitemplate.*;
process filter extends yapiprocess {
    parameter int N_pixels_per_block;
    port yapiininterface-<yapiint>- prod;
    port yapioutinterface-<yapiint>- cons;
    port yapiininterface-<yapiint>- ctrl;

    int p, bc;
    pixel[] blk;
    pixel[] blk_out;
    int filtertype;
}

```



```
filter(String n,int npb,int ft){
    super(n);
    //constructor code omitted
}

public void execute(){
    int enabled;
    ArrayList = new ports ArrayList();
    int[] tokens = new int[2];
    tokens[0] = N_pixels_per_block;
    tokens[1] = 1;
    ports.add(prod);
    ports.add(ctrl);

    while (true){
        enabled = select(ports,tokens);
        /* filter always */
        if (enabled == 0)
            filtertype = ctrl.readint();

        if (filtertype == 1){
            //use algorithm 1
        } else {
            //use algorithm 2
        }
    }
}
```

The user first imports the required library. In order to describe a process, she/he extends the basic class, defines ports and parameters and then overrides the execute method. The filter process behavior has an infinite loop that selects among two inputs: the producer and the controller. If the controller input is selected, the filter type is read. If the controller has not requested any filter type change, the process execution can go ahead and filters the input data.

Finally the top level netlist is described as follows:

```
import metamodel.plt.yapitemplate.*;
public netlist pfcnetlist extends yapinetlist {
    public pfcnetlist(String n){
        super(n);
        producer p = new producer("Prod",48,72);
```

```

consumer c = new consumer("Cons",48,72);
filter f = new filter("Filt",72,1);
controller ctrl = new controller("CTRL",300,150,2);

yapichannel-<yapiint>- cpf = new yapichannel-<yapiint>-("CPF",100);
yapichannel-<yapiint>- cfc = new yapichannel-<yapiint>-("CFC",100);
yapichannel-<yapiint>- ccc = new yapichannel-<yapiint>-("CCC",100);
yapichannel-<yapiint>- cfiltdecision = new yapichannel-<yapiint>-("CFD",100);
yapichannel-<yapiint>- ccp = new yapichannel-<yapiint>-("CCP",100);
//add all componets here
addcomponent(p,this,"ProdInstance");
...
//connect components
connect(p,out,cpf);
...
}
}

```

3.5 Muti-rate Synchronous model of computation

The semantics of this platform can be informally defined as a sequence of rounds. Let $\{P_i\}$ be a set of communicating processes. When a process P_i executes, it execute a sequence of three actions $write_i$, $read_i$ and $execute_i$. $write_i$ writes the content of the internal output buffer O_i to the output ports. $read_i$ read from the input channels and store the data into an internal buffer I_i . Also a process is characterized by two numbers: $P_i.r$ which is the execution rate and $P_i.p$ which is the process priority.

Each round has satisfy the following constraints:

- process P_i has to be executed $P_i.r$ times;
- if $P_i.p > P_j.p$ then P_i has to finish his execution in this round before P_j can execute.

This condition is clearly an ordering of the processes execution depending on rates and priorities. A quantity manager is then needed to model this situation. A detailed description of this platform can be found in the user manual under *metro/doc/polychronyscaled*.

Each process is connected to a quantity manager that is instantiated into a scheduling netlist. The quantity manager is defined as follows:

```
quantity SynchScheduler implements QuantityManager {
    port StateMediumSched[] synchprocesses;
    public SynchScheduler(String n,int st,int nproc){
        super(n);
        //constructor code
        ...}

    public eval void request(event e, RequestClass rc){
        //add requests to the pending list }

    public update void resolve(){
        if ((!_inround) && (_pending.size() == _numberofsynchprocesses)){
            _inround = true;
            toWerePending();
            fairScheduling();
        }else
        if (_pending.size() == _numberofsynchprocesses){
            if (_st == 0){
                fairScheduling();
            };
        }else {
            SynchRequest sr;
            for(int i=0;i<_werepending.size();i++){
                sr = (SynchRequest)_werepending.get(i);
                _notdoevents.add((Object)sr.clone());
            }
        }

    void fairScheduling(){
        int currentpriority,subround;
        _inround = false;
        SynchRequest sr;
        currentpriority = findMaxPriority(_werepending);
        for(int i=0;i<_werepending.size();i++){
            sr = (SynchRequest)_werepending.get(i);
            subround = sr.getClock();
            if ((subround >0) && (sr.getP() == currentpriority) ){
                _inround = true;
                sr.setClock(subround - 1);
                _doevents.add((Object)sr.clone());
            }else{
                _notdoevents.add((Object)sr.clone());
            };
        };
    }
    ...
}
```

3. DEVELOPING PLATFORMS FOR FUNCTIONAL DESCRIPTION

```
public update void postcond(){
    int i;
    int id;
    SynchRequest sr;
    event e;
    for(i = 0; i < _doevents.size();i++){
        sr = (SynchRequest)_doevents.get(i);
        e = sr.getEvent();
        synchprocesses[sr.getId()].setMustDo(e);
        id = sr.getId();
    };
    for(i = 0; i < _notdoevents.size();i++){
        sr = (SynchRequest)_notdoevents.get(i);
        e = sr.getEvent();
        synchprocesses[sr.getId()].setMustNotDo(e);
    };
    _doevents.clear();
    _notdoevents.clear();
    _pending.clear();
}

public eval boolean stable(){
    return true;
}
}
```

If the resolve method is called then the quantity manager checks whether we are in the middle of a round or if a new round should start.

We are in a round, then the *fairScheduling* algorithm is called. This algorithm first compute the maximum priority among all processes with pending requests. Then collects all processes with this priority in a queue. These processes are the one that have to be executed while all the others have to wait. Finally in the *postcond* function orders to execution or not the requested events are given to the processes.

If we are not in a round, then we first wait until all processes have made a request, then we save all requests in an internal array (function *toWerePending*) and then we call the scheduling algorithm.

Chapter Four

Developing platforms for architectural description

There is no significant difference between function and architecture platforms. A platform is a set of library components and a set of rules that define their legal compositions. In the case of architecture it is important to have a way of associating costs and performances to each component and a way of estimating costs and performances of a composition of components. This is very important in order to be able to explore the implementation space and compare solutions to decide which one is the best.

There is though a radical difference between a functional description and an architectural description. Once the denotational description of a function has been mapped onto a platform representing a model of computation, the resulting interconnection of components only implements the original specification. In the case of architecture, an interconnection of library components results in a structure that can implement a set of functions. For instance, a single processor architecture can be used to implement an MPEG decoder or a finite impulse response filter depending on the software that is poured into the program memory.

A platform provider should expose all possible functions that an architecture can support and let the system developer chooses one by determining the architecture behavior in the mapping phase.

4.1 Architecture platforms use-case

We have to consider two different views: platform developers that assemble a platform starting from a set of preexisting library components and platform users that use a platform, configure it and map a set of applications on it. Looking at figure A.1, platform developers would like to define the red point in the platform space starting from the set of IPs that are available in the library. Platform users, instead, want to have control over the total power that the platform can offer which is represented by the light red triangle in the common semantic domain.

Platform developers view. A platform developer has two main concerns:

- she/he wants to define a point in the platform space with certain characteristics (that are usually requested by platform creators) like number of operations per seconds, number of processors, energy per operation etc.
- she/he needs a way of exposing all implementable functions to the platform users. In fact, she/he needs a way of defining the pink triangle and make it available for mapping a function onto the architecture.

Each component in the library should be characterized with parameters that configure its properties. Parameters like bus arbitration policy, instruction set architecture or packet length for a serialization block are examples of configuration options.

Each component should hide its implementation details and only expose its interfaces which declare the services that the component implements and requires. A bus, for instance, exposes the master interface and has ports whose type is a slave interface. Interfaces constraint the way in which components can be connected. A platform developer would connect components together being aware of the valid compositions that can be statically checked by the Metropolis compiler.

Each component is also characterized by costs and performances. These quantities are modeled using quantity managers. When multiple components are connected together quantities could become dependent from each other. Physical time, for instance, is a global quantity that should be the same for all components. Even for quantity that are not global, a

dependence could arise from the composition. Consider a real time operating system and a bus arbiter. When a task mapped on a CPU reads a variable from the memory, the CPU has to ask permission to the bus arbiter. If the bus is busy in another transfer, the opinion of whether the task should continue its execution or not is different for the two schedulers. A coordination mechanism is then needed in presence of multiple quantity managers.

Coordination among quantity managers could become very cumbersome for a platform developer. Ideally, she/he wants the coordination to be totally transparent. A standard coordination algorithm should be provided to the platform developer with very few parameters to act on. For instance, each quantity manager could have a configuration parameter indicating its priority with respect to the other managers. When a platform is built, platform developers could set the priority in order to make the decision of a quantity manager dominate the one of a lower priority manager. In our example, the bus arbiter would have a higher priority than the real time operating system.

Platform user view. A platform user starts with a functional description which is then mapped onto several architectures. Each mapping explores the cost/performance trade-off of the particular architecture under consideration. The goal of this design step is to select the best mapping, or better, the best implementation of the original function at a lower level of abstraction.

A platform user operates in the common semantic domain. The architecture is presented as a set of services and all their valid compositions. The function is described using the same kind of services. The mapping is then the solution of a covering problem.

Two are the features that a user requires:

- a way of modifying the characteristics of an architecture. This requires that a platform exposes a set of parameters through which it is possible to configure its elements. A platform developer has to export the important parameters, characterizing each architecture component, to the level of the platform users.
- A way of “intersecting” the function with the architecture in the common semantic domain. This step can be viewed as a determination of the architecture non-determinism and will be clarified in

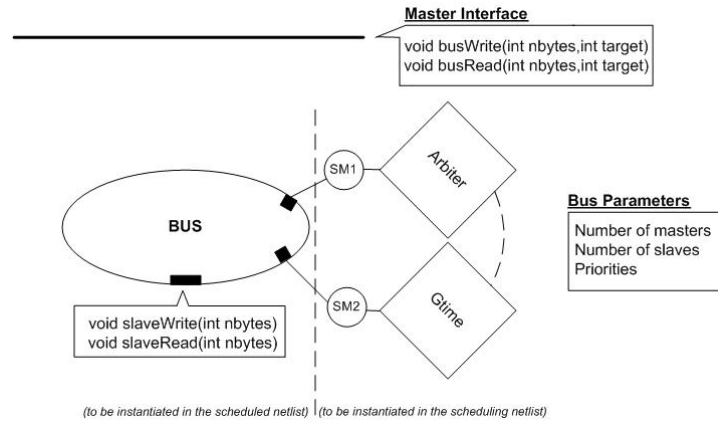


Figure 4.1: A model of a bus

chapter 5. Intuitively, it is a way of assigning pieces of functions to architectural resources that can implement them.

4.2 Architecture platforms development

A platform developer could miss components that she/he needs to build a platform. In this section we provide guidelines for doing two things: providing components and providing an interconnection of components that can implement several functions.

A component offers services that can be used at a cost. A very abstract view of a component is then a medium which exposes some interfaces. Later on a medium can be refined into a netlist if a more detailed description is needed.

Figure 4.1 shows an example of a bus component. The elements to define are:

- the master interface which declares the services that a master can use. In this example there is one service for writing to a specific target and a symmetric service for reading.
- the set of parameters to configure the bus. In this example they are the number of connected masters, the number of connected slaves and the priority of each master.

- A quantity manager that schedules accesses to the bus. This component represents the bus arbiter and decides the ordering of bus accesses among all requests coming from the masters.
- A global time quantity manager. This component is already provided by the Metropolis framework. It is a global quantity and it's in the picture only to show that there is a connection between the bus and the time manager. By using the global time, it is possible to model performances in terms of time required for each operation.
- State media for bus to quantity managers communication. This component implements the communication protocol between the resource and the algorithm that handles it. At sufficiently high level of abstraction, this component implements the identity function passing requests from one component to the other. At lower level of abstraction the communication protocol could be more complicated and the state media have to be refined.

The left hand side of the diagram will be part of a scheduled netlist while the right hand side will be instantiated in the scheduling netlist. The developer should be relieved from the burden of instantiating components, quantity managers and state media and connecting them together. For each component, a function should be provided with the following signature:

```
public myBus addmyBus(netlist theScheduledNetlist,
                     netlist theSchedulingNetlist,
                     ArrayList parameters, int priority)
```

The function will create instances of all components and set their parameters. Then, it will add components to the proper netlist and make connections. The last parameter is needed in the case of multiple quantity managers interacting with each other. The priority parameter will be used in the resolution phase to decide which manager has precedence over the others.

Each component can be described following this basic structure and then refined into a more detailed entity. Refinement could require to move part of the scheduling algorithm from the quantity managers to the scheduled netlist.

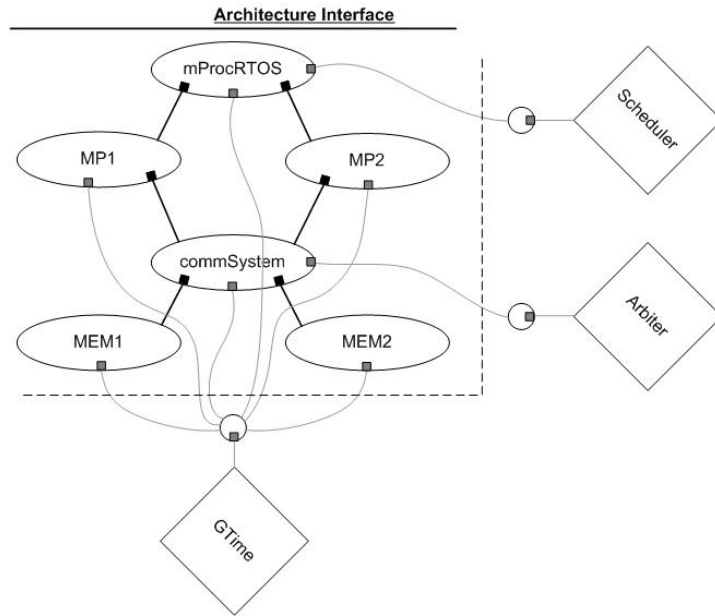


Figure 4.2: Block diagram of a double processor architecture

Consider now a library of components. As platform developers, we want to interconnect them to provide an architecture that can support a variety of applications. We will not deal with the characterization of an architecture as a stand-alone object regardless of the function that is mapped onto it. We rather speculate on how elements of an architectures are connected together, how their parameters are exported and how its potentials are exposed to the platform user.

We chose to provide a double processor architecture. The structure is clear and is shown in figure 4.2.

The architecture exposes an interface which is the real time operating system view of the underlying platform. We select a multi-processor real time operating system, two processors, a communication system, and two memories. We have not specified what kind of components we want to use nor what are the event traces that can be generated by this interconnection of components. Actually, in this abstract architecture there are no processes and hence no events can be generated, or better in the TSM terminology, the only possible execution is the empty behavior.

Before discussing how events are generated and exposed, we describe

how components are interconnected and how parameters are exported.

```
public netlist doubleProc {
    public doubleProc(String n,
                      String comm,
                      String rtossched,
                      double timeslice, double clk1, double clk2){

        //Instance of components
        Netlist dpscheduled = new Netlist(n+'scheduled');
        SchedulingNetlist dpscheduling = new
            SchedulingNetlist(n+'scheduling',true);

        MPRTOS myrtos = new MPRTOS(n,dpscheduled,dpscheduling,
                                   rtossched,timeslice);
        MProc mp1 = MProc(n,dpscheduled,dpscheduling,clk1 );
        MProc mp2 = MProc(n,dpscheduled,dpscheduling,clk2 );
        if (comm == ``xbar``){
            XBar cmm = new XBar(n);
        } else {
            FCFSBus cmm = new FCFSBus(n);
        };
        ....
        connect(myrtos,p1,mp1);
        ...
    }
}
```

The netlist is an example of how the architecture is described. The important thing to notice is that the netlist represents a family of architectures that are not only different because of the parameters of each component but also because the user can decide which communication architecture she/he wants to use. The architecture will still be double-processor-double-memory but it is possible to select a different way of communicating between the two masters and the two slaves. Parameters are simply exported from the components to the top netlist constructor.

We now describe how all the functions that an architecture can implement are exposed to the platform users. Until now, an architecture appears as an interface that offers services. In our example, the multi-processor real time operating system offers services like request from a task to use the CPU, sleep, request of ownership of a semaphore etc. We want to give a description of all legal sequences of those services. We use a non-deterministic requests of the services by a set of tasks. We might




Figure 4.3: Connection of tasks to the architecture

want to limit not only the number of tasks that can be mapped on the real time operating system but also the way in which the services are called or the parameters that are passed to the services. For instance, we cannot read a memory location which is out of the memory space or we cannot execute a bus read if we don't ask the ownership of the CPU before.

Figure 4.3 shows how processes are connected to the architecture model. Before giving a pseudo code for the thread of a process we observe that all quantity managers have to be connected to the process in order to schedule them.

A thread of a process T_i looks like the following:

```
void thread(){
  while(true){
    await{
      (true; ;){ cpuRequest(nondeterminism(int)); } //number of cycles
      (true; ;){ dataRead(nondeterminism(int),nondeterminism(int)); }
                      //address          , number of bytes
      (true; ;){ dataRead(nondeterminism(int),nondeterminism(int)); }
      ...
    }
  }
}
```

```
}
}
```

It is an infinite loop with only one await statement belonging to it. The await has parallel critical sections all enabled at the same time. It semantically says that at each time the while loop is executed, one of the critical section is non-deterministically selected. The parameters of each service are non-deterministic values even if bounds should be given in order to satisfy constraints like memory limitations.

4.3 The single-processor-single-memory architecture

In this section we will create a single processor architecture. We start from the interfaces and then we describe the implementation.

Let's start from the description of a bus. First of all it defines the interfaces to use it. Each bus has its own set of interfaces. The AMBA [1] bus for instance defines the signals that a master and a slave interfaces have to have. Those are very low level characterization of the interconnection and here we want to focus more on the services that are provided/requested by/from a bus. For instance, the Open Core Protocol (OCP) [9] defines different kind of transactions that can be requested by a master and defines also the set of services that a slave has to provide.

We want to provide very basic services to transfer n bytes to a target slave T :

```
public interface MasterInterface extends Port {
    eval void read (int target, int addr, int n, int p);
    update void write(int target, int addr, int n, int p);
}

public interface SlaveInterface extends Port {
    eval void read (int target, int addr, int n);
    update void write(int target, int addr, int n);
}
```

Parameters of this functions are:

- *target*: the target slave;

- *addr*: the address on the target memory space;
- *n*: the number of bytes to transfer;
- *p*: the priority assigned to the master.

Note that we are not concerned with the actual transfer of data even if adding this information is easy. The parameters are needed only to estimate the time required for transferring the *n* bytes. A bus needs two quantity managers: the global time manager that annotates time for the begin and end events of a transaction, and a quantity manager that decides the bus owner (the bus arbiter). When a write or read transaction is requested by a master, a request is issued to the bus arbiter that checks if whether the bus is available. If multiple request are pending, the highest priority master will get the bus ownership. The following piece of code is an idea of how the read service is implemented:

```
public eval      void read (int target, int addr, int n, int p) {
    event e, r;
    br1{@;
    {$
        beg{
            e = beg(getthread(), this.bri);
            _src.setSchedReqClass(e, REQUEST);
            _portSM.request(e, _src);
        }

        $}@};

    {
        _portSlaves[target].busSlaveRead(target, addr, n);
    }
    br2{@;
    {$
        end{
            r = end(getthread(), this.br2);
            _src.setSchedReqClass(r, e , RELEASE);
            _portSM.request(r, _src);
        }
        $}@};
}
```

The service uses two labels: *br1* marks the initial request of the bus while *br2* marks the release of the bus indicating that the tranfer has been suc-

cesfully done. With reference to figure 4.1, when a master calls a read service on a bus a request is issued at the begin event of label *br1*. *_portSM* is a port to the statemedium *S1* that passes the request to the bus arbiter. The arbiter decides if the request can be satisfied in which case the process can continue with its execution and the read on the target slave is called by the bus. When the read finishes, a request to release the bus is issued and the arbiter selects a new owner of the bus (if there are pending requests).

The bus arbiter can also ask the global time manager to account for arbitration overhead in which case the begin event of the read operation on the target slave can begin only if the amount of time needed for arbitration has elapsed. Also the slave read can take a certain amount of time meaning that the end event of the read function can be executed only if such amount of time has elapsed. This algorithm is implemented in the global time quantity manager.

A master could be for instance a CPU. In order for the CPU to be connected to the bus, it must have a port of type *MasterInterface*. A real time operating system is connected on top of the CPU to share the resources among multiple tasks. Instead of going into the details of the CPU code, whose implementation in MMM can be intuitively understood looking at the bus implementation, we want to give an idea of how multiple quantity managers coordinate to guide the architecture execution.

As in the case of the bus and bus arbiter, the real time operating system has a scheduler which is a quantity manager. The question is: how do we coordinate them in order to make the right decision? Each scheduler implements three functions: *resolve*, *postcond* and *stable*. The resolve function looks at the pending requests and based on its scheduling algorithm annotates tags to events. It also decides if an event can be executed or not. The postcond function is used to let each quantity manager know about decisions of the others. The stable function returns a boolean value that should be true only if the schedulers decisions has not changed since last iteration.

The three functions are implemented by the user. The quantity managers are instantiated in a scheduling netlist which is a regular netlist with a resolve method that has to be implemented by the user. The resolve method is called whenever a scheduling decision is needed (for a detailed explanation of this method please refer to [2]). The scheduling

netlist resolve method can recursively call the quantity managers resolve methods until all of them are stable. Finally it calls the postcond methods.

In our case, we could call the bus resolve method first and then the real time operating system scheduler, thus giving the bus arbiter a higher priority. The real time operating system will execute his scheduling algorithm considering only the pending requests that can be executed, without considering tasks that have requested the bus and have to wait until it is available.

After all components are connected together, the real time operating system exposes the following set of services:

```
public interface SwTaskService extends Port {
    eval void request(int n);
    eval void read (int target, int addr, int n);
    update void write(int target, int addr, int n);
    eval void readProtect (int target, int addr, int n);
    update void writeProtect(int target, int addr, int n);
    eval void readLongProtect (int target, int addr, int n, int[] data);
    update void writeLongProtect(int target, int addr, int n, int[] data);
}
```

The *request* method only asks the CPU for computation. Three read and write services are provided. The protect version will first try to acquire a semaphore and then accesses the memory location at address *addr* on target *target*. The long version will transfer the actual data.

The last component that we need are the tasks that use the services provided by the operating system. In our experiment we want to provide services to transfer data and a service for carrying out computation. The task will look like the following:

```
public process SwTask {
    Nondet _target, _addr, _nTimes;
    // Constructor
    port SwTaskService rtos;
    public SwTask(String n, int num_sms, int num_cpus){
        super(n);
        rtos = new SwTaskService[num_cpus];
        ...
    }
}
```



```

public void thread() {
    while(true){
        await{
            (true; ; ) query_space();
            (true; ; ) guard_query_space();
            (true; ; ) claim_space();
            (true; ; ) release_space();

            (true; ; ) query_data();
            (true; ; ) guard_query_data();
            (true; ; ) claim_data();
            (true; ; ) release_data();
            ...
        }
    }
}

public void query_space() {
    rtos[0].readProtect(_target.get(), _addr.get(), 1);
}

public void release_space() {
    int[] tokens = new int[1];

    rtos[0].readLongProtect(_target.get(), _addr.get(), 1, tokens);
    tokens[0] = tokens[0] - _nTimes.get();
    rtos[0].request(1);
    rtos[0].writeLongProtect(_target.get(), _addr.get(), 1, tokens);
}
}

```

Part of the code has been omitted. Nondeterministic data structures have been defined to contain the target slave, the address and the number of bytes to transfer (respectively $target$, $addr$, $nTimes$). These variables are non deterministic quantities in the architecture model but they will be assigned to particular values during the mapping phase.

This architecture provides services for transferring data from and to the memory. They are the same set of services that characterize communication in the Task Transactions level [4] platform described in section 1.1.

For instance, the *release_{space}* method will first read the number of tokens in a FIFO. This variable is stored in memory so it requires to read from the memory an actual data. Moreover the read has to be protected

4. DEVELOPING PLATFORMS FOR ARCHITECTURAL DESCRIPTION

to avoid data corruption. Then the number of tokens in the FIFO has to be updated. This operation also requires an arithmetic instruction to be executed on the CPU. The request function takes into account this overhead. Finally the new number of tokens are written back to memory.

Chapter Five

Mapping a function onto an architecture

Traditionally, a mapping has been thought of as the assignment of pieces of function to architectural resources. If the part of the function S is assigned to resource R it means that R can implement S but R might be able to implement other functions as well. Assigning S to it is a way of restricting the set of its behavior. This chapter uses the more general idea of mapping as intersection of function and architecture in a common semantic domain.

Another important aspect of mapping is scheduling. Function and architecture have two different concurrency models. The amount of concurrency on the architecture side is usually determined by the number computational resources that are available (number of processors for instance). A model of computation like KPN has instead a different model of concurrency. If processes don't communicate at all for instance, they all run concurrently. A scheduler then is needed to adapt this difference.

5.1 Mapping functions onto architectures

The mapping phase has two objectives: implementing a function onto an architecture using its services and selecting the “best” architecture for a particular function (or, more generally, a set of functions belonging to the

5. MAPPING A FUNCTION ONTO AN ARCHITECTURE

same application domain). This last objective is usually called *architecture exploration* or *design space exploration*.

A mapping can be done automatically (if a tool exists to do so) or manually. A designer who is purposed to the mapping phase, starts with a function on one side and a set of architectures on the other side. She/he wants to find a “good” mapping. Words like “best” and “good” are quoted because they are uninterpreted by now, but they implicitly assume that mappings can be compared using a metric representing a trade-off between performance and cost.

An important point that has to be considered is the level of abstraction at which the design exploration is carried out. Each level, in fact, corresponds to a design decision, or a set of decisions, that have to be made in order to reach the implementation level. The common semantic domain should be defined in such a way that these decisions can be made. In this section, we assume that such domain has been selected already which means that function and architecture are described in terms of a common set of services. Note that this assumption is not required by the Metropolis framework but it is suggested by the methodology that we advocate.



Figure 5.1: Examples of function and architecture services requests/offers

Figure 5.1 shows two mapping scenarios. The function F is described as a set of trace where each trace is a sequence of services. The architecture A is pictorially represented as a set resources offering services. For instance, case A shows an architecture with two resources R_1 and R_2 .

Each resource offers two services S_1 and S_2 but only one of them can be in execution at a time.

For each service we shall distinguish between its begin and end event. We shall denote with $beg(S_i)$ the begin event of S_i and with $end(S_i)$ its end event.

Using this simple notation, the function execution trace of case A can be denoted with the sequence $\langle beg(S_1), end(S_1), beg(S_2), end(S_2) \dots \rangle$. The architecture execution is described by a trace of pairs of events because there are two resources. Two examples of its partial behaviors are the following:

$$\begin{aligned} &\langle \{beg(S_1), NOP\}, \{end(S_1), NOP\}, \{NOP, beg(S_2)\}, \{NOP, end(S_2)\} \dots \rangle \\ &\langle \{beg(S_2), beg(S_2)\}, \{end(S_2), end(S_2)\}, \dots \rangle \end{aligned}$$

Even without a formal definition of common semantic domain and intersection, the reader can easily recognize that while the first trace is compatible with the function trace, the second is not because service S_1 is never executed. Moreover, the first trace corresponds to the assignment of S_1 to resource R_1 and service S_2 to resource R_2 .

Case B requires some more observations. Function F requires the concurrent execution of service S_1 and S_2 . If there are no constraints on the begin and end events of these two services that they can both be mapped on resource R_1 and executed in whichever order (S_1 first and then S_2 for instance) because the events are in the function are only partially ordered. If we add the constraint that $beg(S_2)$ as to be executed before $end(S_1)$ then the mapping is possible only if a preemptive scheduling is available on R_1 .

If a mapping is not feasible, intersection of function and architecture results in the empty execution.

5.2 Describing mappings using metamodel synchronization constraints

A mapping can be specified by synchronizing the function execution with the architecture execution. First of all we have to define synchronization of events. Two events e_1 and e_2 are synchronized if $e_1 \in v \iff$

$e_2 \in v$ where v is an event vector. The synchronization relation is transitive.

Synchronization among events can be enforced using constraints. Consider for instance case A of figure 5.1. We want to map service S_i on resource R_i . This result can be obtained by synchronizing the begin event of S_i in the function side, with the begin event of S_i on resource R_i in the architecture side. We also want to synchronize their end events.

Synchronization is obtained using LTL formulas whose declaration is as follows:

```
ltl synch(e1,e2,...[: v1@(e1,i)==v2@(e2,i),...])
```

it means that all the events in the first part are in the synchronization relation. The second part is optional and is used to assign values during mapping. In chapter 4 we have seen that values in the architecture are non-deterministic. They become deterministic when during mapping where an actual value is decided for them. For instance, the memory address of a variable is decided in the mapping phase.

The LTL formulas are interpreted and synthesized by the simulator that will try to schedule events in order to satisfy them.

AppendixAppendix

Appendix A

Platform-Based Design example

The platform-based design methodology is a recursive paradigm where the action of mapping a function onto an architecture generates a new function described at a lower level of abstraction and therefore more detailed than the original one.

A design process should start with a denotational description of the function that we want to implement, plus a set of constraints that the implementation has to satisfy. Filtering a signal $x(t)$, for instance, can be denotationally described as $x(t) \otimes h(t)$, namely the convolution of the signal with the filter impulse response $h(t)$. Design constraints are usually specified as propositional formulas over the system quantities. In the case of filtering, for example, we can specify a lower bound on the off-band signal attenuation. Constraints specified at this level of abstraction are propagated down to all subsequent levels, until the implementation level is reached.

While constraints are propagated in a top-down fashion, performances are abstracted in a bottom-up manner. Performance abstraction is the process of hiding details that are not relevant for the level of abstraction under consideration. In fact, each level of abstraction focuses on a particular design choice on which only few quantities have impact. Abstraction of quantities that are not relevant is essential for speeding up the design space exploration.

In this section we focus on one step of the design flow characterized by a function, an architecture and the mapping of the former onto the latter. We use a simplified logic synthesis flow as a representative example.

Figure A.1 shows the design process. The function is described in

A. PLATFORM-BASED DESIGN EXAMPLE




Figure A.1: Platform based representation of a simplified logic synthesis flow

the register transfer level (RTL) domain. In this domain a function is described as interconnection of combinational logic blocks communicating through registers. Furthermore, each combinational block takes zero time to compute its logic function.

The platform is composed of all possible logic functions that can be implemented on a chip using standard cells technology. The library of components, from which the platform is built up, contains pre-characterized logic functions that are usually custom designed to achieve extremely high performances. Each library element is characterized by its logic function, performance and cost. An OR logic gate, for instance has a truth table, gate delay and power consumption associated with it.

A common semantic domain for the RTL functional description and the standard cell platform is the domain of all circuits built out of 2-inputs NAND gates. In fact, every logic function can be expressed only using NAND gates. Mapping an RTL function onto the standard cell platform can be done in the following way. First, we analyze the RTL description and, for each combinational block, we express its function using 2-inputs NAND gates. Secondly, we express each library cell logic function using the same elementary gate. Finally, the problem reduces to a minimum cost covering of the original function with the library elements with a given cost function (e.g. power or minimum slack).

The covering problem is an optimization problem subject to the constraints coming from the original specification. Typical constraints that are associated with the functional specification are clock speed, inputs arrival time and output required time. After the covering algorithm has finished, those constraints must be checked (using timing analysis). If they are not met, the designer has two choices: going back and modifying the RTL description by using a different block partitioning, or moving the covering algorithm out of the local minima.

The resulting netlist is an interconnection of standard cells which is indeed a new function F' . This function represents a refinement of the original one which was described at a much higher level of abstraction. F' is the new specification for the next design step whose platform library is composed of transistors and wires. Original constraints are propagated further down and must be satisfied by the transistor level implementation.

The simple design process described above turns out to be very general. Once a common semantic domain has been defined so that both the function and the platform elements can be expressed using a common mathematical formalism, design exploration reduces to a covering problem. The covering algorithm has to minimize a cost function which is the sum of the costs of each platform component that has been used during the covering. The whole optimization problem is subject to the original constraints.

Appendix B

Tagged Signal Model Definitions

The tagged signal model is a denotational framework to compare models of computation. In this framework an event e is an element of the cross product $V \times T$ where V is a set of values and T is a set of tags. An order relation defined on T induces an ordering of events. A signal s is a set of events and can be viewed as subset of $V \times T$. It is useful to define a set of signal s and also the set S^N of all sets of N signals. A behavior of a process is an element of the set S^N and a process P is a subset of S^N . $s \in S^N$ satisfies a process if $s \in P$ and it is called a behavior of the process P .

Bibliography

- [1] ARM. <http://www.arm.com/products/solutions/ambahomepage.html>.
- [2] Felice Balarin and Yosinori Watanabe. Metamodel language.
- [3] Felice Balarin, Yosinori Watanabe, Harry Hsieh, Luciano Lavagno, Claudio Passerone, and Alberto Sangiovanni-Vincentelli. Metropolis: An integrated electronic system design environment. *IEEE Computer*, Apr 2003.
- [4] J-Y. Brunel, E.A. de Kock, W.M. Kruijtzer, H.J.H.N. Kenter, and W.J.M. Smits. Communication refinement in video systems on chip. *7th International Workshop on Hardware/Software Co-Design*, May 1999.
- [5] W. J. M. Smits P. van der Wolf J.-Y. Brunel W. M. Kruijtzer P. Lieverse K. A. Vissers E. A. de Kock, G. Essink. Yapi: Application modeling for signal processing systems. *Proceedings of the Design Automation Conference*, June 2000.
- [6] Gilles Kahn. The semantics of a simple language for parallel programming. *Proceedings of IFIP Congress 74*, August 1974.
- [7] K. Keutzer, S. Malik, A. R. Newton, J. Rabaey, and A. Sangiovanni-Vincentelli. System level design: Orthogonalization of concerns and platform-based design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 19(12), December 2000.
- [8] E. Lee and A. Sangiovanni-Vincentelli. A denotational framework for comparing models of computation. *Technical Memorandum UCB/ERL M97/11*, November 1997.
- [9] Open Cores Protocol. <http://www.ocp.org>.