


UNIVERSITY OF
NEWCASTLE



University of Newcastle upon Tyne

COMPUTING SCIENCE

Incremental Development of a Distributed Real-Time Model of a Cardiac
Pacing System using VDM

H. D. Macedo, P. G. Larsen and J. Fitzgerald

TECHNICAL REPORT SERIES

No. CS-TR-1059 November, 2007

Incremental Development of a Distributed Real-Time Model of a Cardiac Pacing System using VDM

Hugo D. Macedo, Peter G. Larsen and John Fitzgerald

Abstract

The construction of formal models of real-time distributed systems is a considerable practical challenge. We propose and illustrate a pragmatic incremental approach in which detail is progressively added to abstract system-level specifications of functional and timing properties via intermediate models that express system architecture, concurrency and timing behaviour. The approach is illustrated by developing a new formal model of the cardiac pacemaker system proposed as a “grand challenge” problem in 2007. The models are expressed using the Vienna Development Method (VDM) and are validated primarily by scenario-based tests, including the analysis of timed traces. We argue that the insight gained using this staged modelling approach will be valuable in the subsequent development of implementations, and in detecting potential bottlenecks within suggested implementation architectures.

Bibliographical details

MACEDO, H. D., LARSEN, P. G., FITZGERALD, J..

Incremental Development of a Distributed Real-Time Model of a Cardiac Pacing System using VDM
[By] H. D. Macedo, P. G. Larsen, J. Fitzgerald.

Newcastle upon Tyne: University of Newcastle upon Tyne: Computing Science, 2007.

(University of Newcastle upon Tyne, Computing Science, Technical Report Series, No. CS-TR-1059

Added entries

UNIVERSITY OF NEWCASTLE UPON TYNE
Computing Science. Technical Report Series. CS-TR-1059

Abstract

The construction of formal models of real-time distributed systems is a considerable practical challenge. We propose and illustrate a pragmatic incremental approach in which detail is progressively added to abstract system-level specifications of functional and timing properties via intermediate models that express system architecture, concurrency and timing behaviour. The approach is illustrated by developing a new formal model of the cardiac pacemaker system proposed as a “grand challenge” problem in 2007. The models are expressed using the Vienna Development Method (VDM) and are validated primarily by scenario-based tests, including the analysis of timed traces. We argue that the insight gained using this staged modelling approach will be valuable in the subsequent development of implementations, and in detecting potential bottlenecks within suggested implementation architectures.

About the author

Hugo Macedo is a Doctoral Student in Computer Science at Minho University. His research interests are the use of formal methods, requirements engineering, model driven software and education. After he finished his degree on Maths and Computer Science he joined the Overture Community (www.overturetool.org) and has been a working on VDM since then. Much of the work reported here was undertaken while he was visiting the Engineering College of Aarhus.

Peter Gorm Larsen is Professor of Computer Technology and Embedded Systems at The Engineering College of Aarhus, Denmark and an independent consultant. An authority on system modelling, particularly the Vienna Development Method, he has pioneered the development of industrial-strength tool support for model-oriented specification languages, heading the group that initially developed VDMTools(R) now owned by CSK.

John Fitzgerald is Reader in Computing Science at Newcastle University. His research concerns the use of formal methods to support the design of systems that reconfigure in response to threats. He is co-Investigator in the Trustworthy Ambient Systems project and is Chairman of Formal Methods Europe.

Suggested keywords

FORMAL METHODS,
REAL-TIME SYSTEMS,
GRAND CHALLENGE,
VDM

Incremental Development of a Distributed Real-Time Model of a Cardiac Pacing System using VDM

Hugo Daniel Macedo¹, Peter Gorm Larsen², and John Fitzgerald³

¹ Minho University, Portugal

² Engineering College of Aarhus, Denmark

³ School of Computing Science, Newcastle University, UK

hmacedo@di.uminho.pt, pgl@iha.dk, John.Fitzgerald@ncl.ac.uk

Abstract. The construction of formal models of real-time distributed systems is a considerable practical challenge. We propose and illustrate a pragmatic incremental approach in which detail is progressively added to abstract system-level specifications of functional and timing properties via intermediate models that express system architecture, concurrency and timing behaviour. The approach is illustrated by developing a new formal model of the cardiac pacemaker system proposed as a “grand challenge” problem in 2007. The models are expressed using the Vienna Development Method (VDM) and are validated primarily by scenario-based tests, including the analysis of timed traces. We argue that the insight gained using this staged modelling approach will be valuable in the subsequent development of implementations, and in detecting potential bottlenecks within suggested implementation architectures.

1 Introduction

Formal models have a valuable role to play in validating requirements and designs for real-time distributed systems in early development stages. Rapid feedback from the analysis of such models has the potential to reduce the risk of expensive re-working as a consequence of the late-stage detection of defects. However, models that incorporate the description of functionality alongside timing behaviour and distribution across shared computing resources are themselves potentially complex. Moving too rapidly to such a complex model can increase modelling and design costs in the long run. In order to gain full value from formal modelling and analysis, a systematic approach to constructing and validating models is required.


Our current work is focussed on the development and industrial application of formal modelling techniques that satisfy the requirements discussed above. We have developed and applied technology based on the Vienna Development Method (VDM) [1,2,3] and its tool support (VDMTools [4]). Recent work has developed modelling abstractions and test-based analysis tools that support the object-oriented description of distributed real-time systems [5,6]. Our experience applying formal modelling techniques in a variety of industry sectors suggests that an approach to modelling such distributed real-time systems should permit the staged and controlled construction of a formal model, with opportunities for validation at each stage. We have proposed such a staged approach as part of the methodological guidelines accompanying the VDMTools [7].

This paper reports a study in which we have assessed the feasibility of applying an incremental approach to model construction by developing a model for an artificial cardiac pacemaker [8]. The pacemaker specification has been offered as the basis of a “grand challenge” problems in computing [9] by the Software Quality Research Laboratory at McMaster University. An important characteristic of the challenge is that it includes system-level requirements affecting hardware as well as software. This paper demonstrates how such cross-disciplinary requirements can be introduced gradually into a model in a phased fashion, along with validation of the functional and timing properties expressed in the model.

We introduce the Pacemaker system in Section 2. Section 3 briefly introduces the VDM technology used, our phased approach to model construction and the tool support. The progressive development of the Pacemaker model is described, illustrated by extracts from the series of VDM models developed (Section 4). The test-based approach to validation is discussed in Section 5, including how timing conditions can be checked using this technology. Finally Sections 6 and 7 discuss related work and draw conclusions from the study.

2 The Pacemaker System and Environment

In this study, the pacemaker is treated as an embedded system operating in an environment containing the heart. We first review the elements of the environment that interact with the pacemaker (Section 2.1) and then consider the elements of the pacemaker system itself (Section 2.2).



2.1 Environment: The Heart

The human heart serves as a pump for the circulatory system. It is a muscular shell around four chambers (called atria and ventricles) which contract and relax periodically under the control of natural electrical stimuli. A natural pacemaker orchestrates the functioning of the pump, discharging electrical pulses at specific points (see Fig. 1). In normal functioning, a discharge is made at the sinus node; the discharge subsequently reaches the atrioventricular (AV) node which amplifies it, stimulating the ventricles. If the natural pacemaker is malfunctioning, a physical condition termed Bradycardia may arise in which the heart rate falls below the level expected for the patient. To normalise the heart rate, an artificial pacemaker may be implanted to aid or replace the natural pacemaker. Physicians measure the heart’s performance using, among other parameters, the bpm (beats per minute) rate of the heart. We use the term *pulse* and *pulses per minute* in reference to pacemaker activity, whereas *beat* and *beats per minute* refer to heart activity.

Fig. 1. The natural pacemaker

2.2 System: Artificial Pacemaker

An artificial pacemaker (referred to subsequently as a *pacemaker*) is a system composed of:

Leads: One or more wires, normally two, that both sense and discharge electric pulses.

Device: The implanted batteries and controller.

Device Controller-Monitor (DCM): An external unit that interacts with the device using a wireless connection (not modelled in this paper.)

Accelerometer: A unit inside the device measuring body motion in order to allow modulated pacing.

A typical configuration consists of one lead attached to the right atrium and another to the right ventricle. The pacemaker has several operating modes that address different malfunctions of the natural pacemaker. The specification document [8] identifies 18 operating modes controlling 26 variables and each of the variables can be configured within a value range. Most of the variables are time-related parameters, defining such properties as the interval between a pace in the atrium and the ventricle or the number of pulses per minute the device should deliver to a given chamber.

The operating modes of the device are classified using a code consisting of three or four characters. For the examples in this paper, the code elements are: chamber(s) paced (“O” for none, “A” for atrium, “V” for ventricle, “D” for both), chamber(s) sensed (same codes), response to sensing (“O” for none in this paper), response to sensing (“O” for none in this paper) and a final optional “R” to indicate the presence of rate modulation in response to the physical activity of the patient as measured by the accelerometer. “X” is a wildcard used to denote any letter (i.e. “O”, “A”, “V” or “D”). Thus “DOO” is an operating mode in which both chambers are paced but no chambers are sensed, and “XXXXR” denotes all modes with rate modulation.

3 VDM Modelling Technology for Distributed Real-Time Systems

In our modelling work, we have used VDM [1]. Three dialects of the VDM modelling language are in use, each supporting different forms of system specification. VDM-SL [10] provides facilities for the functional specification of sequential systems with basic support for modular structuring. VDM++ [3] extends VDM-SL with features for object-oriented modelling and concurrency. VICE (VDM++ In Constrained Environments) further extends VDM++ with features for describing real-time computations [11] and distributed systems [5]. Each dialect has formally defined syntax, static and dynamic semantics which extend those of the ISO Standard VDM-SL language [12]. For a detailed introduction to VDM++, the reader is referred to the texts and the VDM Portal [13]. In the remainder of this section, we focus on the features of VDM++ and VICE that have a major role in the modelling of distributed real-time systems.

3.1 Basic VDM Notations

A model in VDM-SL, is composed of type definitions built from simple abstract types such as *bool* or *nat*, and constructors such as sequences and records. Types may be restricted by predicate invariants. Persistent state variables may be defined. Operations that may modify the state can be defined implicitly, using pre- and postcondition expressions, or explicitly, using imperative statements. Functions are similar to operations except they may not refer to state variables, and are side-effect free.

An object-oriented model in VDM++ is composed of class specifications similar in many ways to VDM-SL models and which may use single or multiple inheritance. The state of an object consists of typed *instance variables*. Operations in VDM++ are re-entrant and their invocation is defined with synchronous (rendezvous) semantics. Operation execution may be constrained by specifying a permission predicate [14], a Boolean expression over *history counters* that acts as a guard for the operation, for example to express mutual exclusion. History counters are maintained per object to count the number of requests, activations and completions per operation.

VDM++ classes may be active or passive. Active classes represent entities that have their own thread of control; passive classes are always manipulated from the thread of control of another active class. A thread is a sequence of statements that is executed to completion, at which point the thread dies. The thread is created whenever the object is created but the thread needs to be started explicitly using a *start* operator. It is possible to specify threads that do not terminate.

Extensions to VDM++ (VICE) support the description and analysis of real-time embedded and distributed systems [15,16]. These include primitives for modelling deployment to a distributed hardware architecture and support for asynchronous communication. Two predefined classes, *BUS* and *CPU*, are available to the specifier to construct the distributed architecture in the model. User-defined classes can be instantiated and deployed on specific *CPUs*. The communication topology between the computation resources in the model can be described using the *BUS* class.

The semantics of VDM++ is extended with a notion of time such that any thread that is running on a computation resource or any message that is in transit on a communication resource can cause time to elapse. Models that contain only one computation resource are compatible to models in plain VDM++.

Operations may be specified as asynchronous, allowing the caller to resume its own thread of control after the call is initiated. A new thread is created, scheduled and started to execute the body of the asynchronous operation. Statements (*duration* and *cycles*) may be used in operation bodies to specify time delays that are, respectively independent of or dependent upon processor capacity. The time delay incurred by the message transfer over the *BUS* can be made dependent on the size of the message being transferred.

3.2 An Incremental Approach to Model Construction

Faced with the challenge of developing VDM++ models of distributed real-time systems, we have proposed a staged approach [7] which reflects the capabilities of each of the VDM modelling languages.

The analysis of informally expressed requirements leads to a first abstract model giving system-level specification of behaviour. The basic VDM-SL language is well suited to this level of description. Based on the abstract model, we introduce a static architecture, creating a sequential (non-concurrent) model with structure expressed using the features of VDM++. This model would then be extended to become a concurrent VDM++ design model. The concurrent design model itself is then extended with real-time information using the VICE extensions, and additionally distribution over processors can be described also using the VICE extensions. At this stage it may prove necessary to revisit the concurrent design model, since design decisions made at that stage may prove to be infeasible when real-time information is added to the model (for instance, the model may not be able to meet its deadlines).

The ability to validate the intermediate models developed in this process makes it possible to identify requirements and design defects at an early stage. The initial abstract model need not be directly executable, but subsequent models are likely to be so, making it possible to conduct extensive tests in order to validate design decisions. The VDMTools are intended to provide extensive support for scenario-based testing as a form of validation.

We do not claim that the models introduced at each stage in our approach are formal refinements of their predecessors, although this may sometimes be the case. Our intended output is a comprehensive model of the target system that can serve as a basis for subsequent development, possibly using refinement. We are therefore introducing detail in a staged manner, where the executions at each level might, informally, be seen as providing a finer level of granularity than its predecessor.

3.3 VDM Tool Support

VDM is supported by an industry-strength tool set, VDMTools, owned and developed by CSK Systems [17]. VDM and VDMTools have been used successfully in several large-scale industrial projects, e.g. [18,19]. The tools offer syntax, type and integrity checking capabilities, code generators, a pretty printer and an application programmer interface. The main support for model validation is by means of an interpreter allowing the execution of models written in the large executable subset of the language.

Scenarios defined by the user are essentially test cases consisting of scripts invoking the model's functionality. The interpreter executes the script over the model and returns observable results as well as an *execution trace* containing, for each event, a time stamp and an indication of the part of the model in which it appeared. A separate tool (an Eclipse plug-in) called *showtrace* has been developed for reading execution traces, displaying them graphically so that the user can readily inspect behaviour after the execution of a scenario, and thereby gain insight into the ordering and timing of exchange of messages, activation of threads and invocation of operations.

The existing tools have been further extended to allow explicit logical statements of expected system-level timing properties (termed *timing conjectures*) which can be checked against execution traces [6]. Fig. 2 shows the *showtrace* output resulting from the analysis of three validation conjectures (C1-C3) from the pacemaker study (see Section 5). The main window shows a fragment of the execution trace, with time on the horizontal axis. Processing on each architectural unit is shown by horizontal lines (colours

are used to denote thread start-up, kill and scheduling). Thin arrows indicate message passing and fat arrows indicate thread swapping. The conjectures are shown at the bottom of the window. Circular marks on the traces show conjecture violations, e.g. the circle showing a counterexample to the timing conjecture C3, where an event occurrence breaches an expected temporal separation.

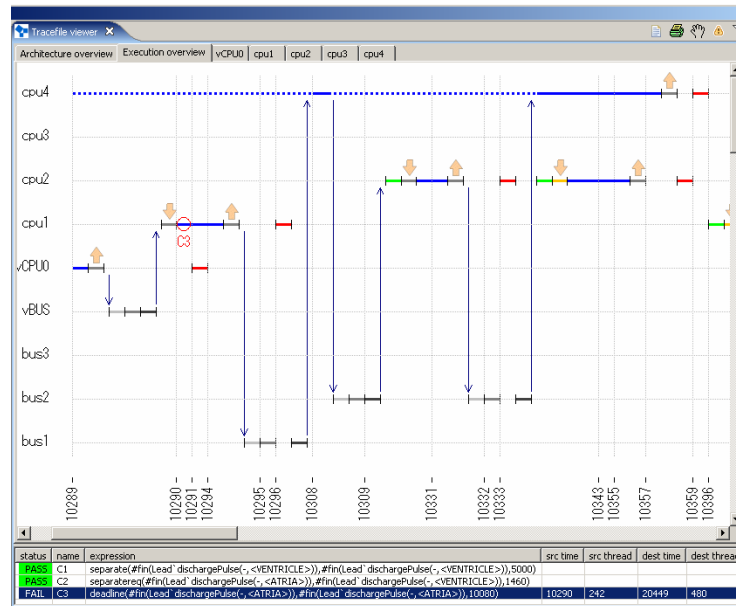


Fig. 2. Showtrace tool demonstrating validation conjecture violation

4 The Pacemaker Models

The overall purpose of the modelling work on the Pacemaker has been to clarify and validate the system’s informally stated requirements [8]. Following our staged approach, in order to manage the complexity of the model itself, the construction was done in four steps, each involving the construction of a new model at a lower level of abstraction from its predecessor. We will term them *Abstract*, *Sequential*, *Concurrent* and *Distributed Real-Time* (DR-T) respectively.

Note that we do not claim that this is a *formal* refinement process. The initial *Abstract* model is informally refined by a *Sequential* model by adding structuring information. Neither of these models the concurrency of the environment and the system; instead they simulate fixed time steps controlled from the environment. In the *Concurrent* model both the environment and the relevant parts of the system are organised with concurrent threads that are synchronized by permission predicates. In each of these three models, time is present explicitly as an abstraction whereas in the final *Distributed*

Real-Time model time is implicit, allowing us to express more realistic timing behaviour while validating this model.

4.1 Abstract Model

The first model is expressed in VDM-SL, the simplest of the VDM modelling languages, lacking the object-orientation and concurrency features of VDM++. The model consists of several modules, each corresponding to an operating mode of the pacemaker. Each module defines a single function called `PacemakerM` (where `M` is the mode), derived from the requirements. Each `PacemakerM` function defines a timed output trace that should result from an input trace of sensed data. Thus we have defined types as follows:

```
SenseTimeline = set of (Sense * Time);  
ReactionTimeline = set of (Reaction * Time);
```

where `Sense` and `Reaction` are enumerated types representing the presence or absence of a pulse.

Each pacemaker function is expressed in VDM-SL in an implicit style by means of a postcondition characterising the events trace that should result from correct functioning of the pacemaker over an input sense trace. The implicit style is used because it is not intended that the function should be directly executed; it serves primarily as a means of clarifying requirements.

```
PacemakerM (inp : SenseTimeline) r : ReactionTimeline  
post ...
```

As an example, consider the specification of the pacemaker in the “DOO” mode. In the requirements document, this mode has a requirement stating that the pacemaker must deliver 60 pulses per minute (PPM) to the atria. The model defines several constants that relate to the Lower Rate Limit (LRL, the number of pace pulses delivered per minute in the absence of sensed activity in an interval starting at a paced event); Upper Rate Limit (URL, maximum rate at which the paced ventricular rate will track sensed atrial events) and the Fixed AV Delay (the programmable period from an atrial event to a ventricular pace in milliseconds). For some such parameters, we have selected nominal values from permitted ranges; the model can be readily extended to accommodate tolerances.

values

```
LRL      : nat := 60;  
URL      : nat := 120;  
FixedAV  : nat := 150;
```

functions

```
PacemakerDOO(inp:SensedTimeline, n:Time) r:ReactionTimeline  
post  
  let nPulsesAtria = card {i | i in set r & i.#1 = <ATRIUM>},  
      nPulsesVentricle = card {i | i in set r & i.#1 = <VENTRICLE>}  
  in nPulsesAtria / n >= (LRL / 60) / 1000
```

```

and
  nPulsesVentricle / n <= (URL / 60) / 1000
and
  forall mk_(<ATRIUM>,ta) in set r &
    (exists mk_(<VENTRICLE>,tv) in set r & tv = ta + FixedAV);

```

The behaviour of the pacing mechanism in this mode is defined by stating that the result should contain the number of pulses required for LRL and URL. The numeric conversions are needed because the timeline unit is milliseconds. Note also that the input trace is not used in the body of the function, reflecting the requirement that this mode issues pacing pulses regardless of any sensed data.

The abstract models support the formalisation of our understanding of the system requirements. For example, while we were modelling this mode we noticed that the requirements in [8] for some modes place constraints on ventricular pace events even when the ventricle is not being paced. We also identified areas of incompleteness, for example the requirements as modelled in `PacemakerDOO` above do not take account of some unstated requirements on the intervals between atrium pulses. A revised post-condition would be as follows:

```

post
  let LRLint = 1000
  in forall mk_(<ATRIUM>,ta) in set r &
    ((exists mk_(<VENTRICLE>,tv) in set r & tv = ta + FixedAV)
     and
     (ta + LRLperiod > n
      or
      (exists mk_(<ATRIUM>,ta3) in set r & ta + LRLint = ta3)))
  and
  forall mk_(<VENTRICLE>,tv) in set r &
    (exists mk_(<ATRIUM>,ta2) in set r & ta2 + FixedAV = tv);

```

Here the intervals between atrium and ventricle pulses are `FixedAV` as before, but we reach the LRL requirement by imposing an fixed interval between atrium paces, in this case 1000. The URL is also trivially satisfied because we only allow ventricle pulses after each atrium pulse, as stated in the right hand side of the final conjunct. This formalisation is more restrictive than our original; indeed, it is one valid implementation of the more abstract initial model. In a full modelling process, we would expect to explore the model further with domain experts, leading to a revised requirements statement.

As a further product of the abstract modelling phase, we can use the postconditions of the `PacemakerM` functions as test oracles on the models developed in the subsequent phases, provided suitable abstraction functions are implemented. We were also able to use this model to help design the validation conjectures that were applied to the analysis of the final (distributed real-time) model.

4.2 Sequential Model

The sequential design model describes both the data that is to be computed, and how it is to be structured into static classes, without commitment to a specific dynamic ar-

chitecture. For the Pacemaker example, a static architecture is shown in Fig. 3 where we can see the system represented by its prime class `Pacemaker` coexisting with its `Environment` in a given `World`. This model is then used to analyse the system behaviour without taking concurrency into account. In the sequential step, the environment affects the flow-of-control by passing signals to the pacemaker system. The environment also provides the simulated time increments using the `Timer` class.

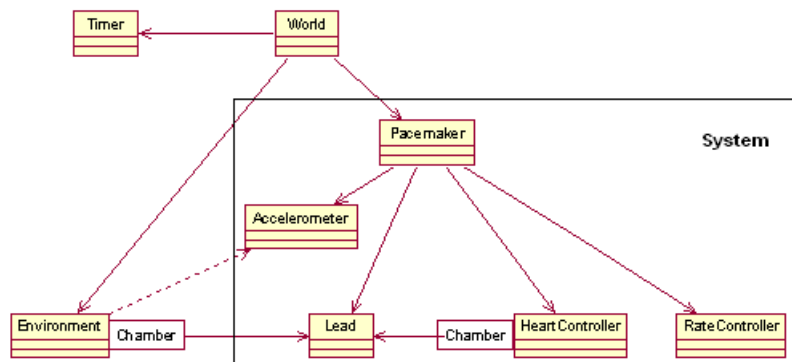


Fig. 3. A UML class diagram describing the static structure

The main feature of the static architecture is its division into environment and pacemaker system. The `Environment` class that controls the production of stimuli to the system model delivered via `Leads`. The `Pacemaker` class represents the technical system, with a `HeartController` that provides the core functionality monitoring incoming stimuli and generating pulses. The `RateController` is used for rate adaptation control in “XXXXR” operating modes, and will not be further discussed in this paper (for full details see [20,21]) and the `Accelerometer` coping with the motion data.

The environment class contains an explicit `Run` operation which operates a form of command loop, stepping through the input timeline of sensed events, delivering signals to the system (through the operation `createSignal`). Having created the signals, it calls a `Step` operation in `rateController` and later in `heartController`, the two parts of the system model that will become concurrent in subsequent models. In addition the global time variable is stepped forward.

```

public
Run: () ==> ()
Run () ==
  (while not (isFinished()))
  do
    (createSignal());
    Pacemaker`rateController.Step();
    Pacemaker`heartController.Step();
  
```

```

        World`timerRef.StepTime();
    );
);

```

On the system side, the `HeartController` paces the heart using operation `pace` which calls mode-dependent operations that determine a response intended to satisfy the postcondition of the corresponding operation in the abstract model.

```

public
Step : () ==> ()
Step () ==
    (pace();
    for all key in set dom leads
    do leads(key).Step();
    );

```

The sequential model is executable. We are able to validate the model coupled with the informal requirements by running a series of tests through the VDM++ interpreter.

4.3 Concurrent Model

In the third phase of model development, concurrency is introduced. However the static structure from Fig. 3 is maintained, with the exception of a more elaborate modelling of time [20]. The environment model is freed of the responsibility to control the system model. The concurrent model has the same structure as the sequential one, but the stepping mechanism is substituted by threading and synchronisation. The previously sequential `Run` operation is now a thread that still controls time, but not the control flow.

```

thread (start(new ClockTick(threadid));
    while true do
        ( if busy
            then createSignal();
            World`timerRef.NotifyAndIncTime();
            World`timerRef.WaitRelative(1);
        );
    );

```

With the introduction of concurrency, the need for synchronisation arises and this is achieved through the use of permission predicates. For instance, the reactions of the system are collected in a variable inside the `Environment` class. There is an operation called `handleEvent` that is used to write values on this instance variable. We need to avoid a race condition caused by the concurrent access to this variable and thus a *mutex* constraint is introduced.

```

sync mutex (handleEvent);

```

The former `Step` operations also become threads. The `HeartController Step` now becomes a similar thread but instead of the sequential model time increment: 200 milliseconds, now the thread sleeps using a Wait/Notify design pattern:

```

thread while true do
    ( World`timerRef.WaitRelative(200);

```

```

        pace();
    );

```

4.4 Distributed Real-Time Model

The final step in model construction is the introduction of distribution over CPUs in a topology determined by the configuration of a bus. Time “annotations” allow time to be built in and managed by the VDM++ interpreter so again the static structure is preserved except that the explicit modelling of time now disappears.

The deadlines and periods in the requirements [8] are stated in milliseconds and in certain cases in fractions of milliseconds. In order to be able to obtain the correct level of granularity microseconds will be used as the unit of time below.

The `Environment` class now only delivers signals to the system as a periodic thread:

```

thread periodic (1000,10,900,0) (createSignal);

```

The operation `createSignal` will be invoked approximately every 1000 time units, each millisecond and a jitter of up to 10 time units can be allowed for the periodic invocation of this thread. The third parameter indicates that there is going to be at least 900 time units between two wake-ups of this periodic thread. Finally, the last parameter indicates that no offset is required for the invocation of it. This is a feature that is most valuable if a number of threads are started at the same time, and there is a desire to carry them out in a special order.

Expressing Time Requirements Time “annotations” are used to record the durations allocated to particular actions. For example, consider the requirement that “The Atrial pulse width must be 0.05 milliseconds.” [8], appendix. This is a requirement on the operation that describes the discharge of a pulse. The `duration` statement of VDM++ is used to specify the width in microseconds.

```

private async
dischargePulse : Pulse * Chamber ==> ()
dischargePulse (p,c) ==
    if(c = <ATRIUM>)
        duration(50) World`env.handleEvent(p,c,time);

```

Notice also the use of `async` so that the caller of this operation will not block waiting for it to terminate.

Modelling the System Distribution The system class defines the distribution architecture. Instance variables are defined as follows (all public static) to define the physical objects related to the problem domain:

```

atriaLead      : Lead          := new Lead(<ATRIUM>);
ventricleLead  : Lead          := new Lead(<VENTRICLE>);
accelerometer  : Accelerometer := new Accelerometer();
rateController : RateController := new RateController();
heartController : HeartController := new HeartController();

```

The architectural components in the model are four CPUs. In the CPU definitions, we indicate the scheduling algorithm and processor capacity (Fixed Priority (FP) and First-Come-First-Served (FCFS) respectively). For the cases of the CPUs that will just run a single thread we used a FCFS scheduling algorithm. For the case of the CPU that will contain different threads (the `rateController` and `heartController` ones) we opted for a fixed priority one, because it is reasonable to assume that the stimulation of the patient heart is more important than adjusting the rate of the stimulation.

```
cpu1 : CPU := new CPU (<FP>, 1E3);
cpu2 : CPU := new CPU (<FCFS>, 1E3);
cpu3 : CPU := new CPU (<FCFS>, 1E3);
cpu4 : CPU := new CPU (<FCFS>, 1E3);
```

In order to define the communication topology, we create three bus objects linking the specified CPUs with a certain bandwidth (1E6) and the chosen network control protocol, in this case FCFS.

```
bus1 : BUS := new BUS (<FCFS>, 1E6, {cpu1, cpu4});
bus2 : BUS := new BUS (<FCFS>, 1E6, {cpu2, cpu4});
bus3 : BUS := new BUS (<FCFS>, 1E6, {cpu3, cpu4});
```

The final element in this system class is the constructor operation that deploys the functionality across the resources. Here the leads are deployed on two processors (representing the physical wires and lead controllers) and the accelerometer is deployed using the same approach. The remaining devices are deployed on `cpu4`:

```
public Pacemaker: () ==> Pacemaker
Pacemaker () ==
  (cpu1.deploy (atriaLead);
   cpu2.deploy (ventricleLead);
   cpu3.deploy (accelerometer);
   cpu4.deploy (rateController);
   cpu4.deploy (heartController);
   cpu4.setPriority (HeartController `pace, 3);
   cpu4.setPriority (RateController `adjustRate, 1);
  );
```

Implicitly there is always a virtual CPU and BUS where elements that are not deployed to a explicit CPU run. Also the communication between objects in different CPUs with no explicit bus connecting them will occur using the virtual bus.

5 Validation

5.1 Validation of Abstract, Sequential and Concurrent Models

A systematic testing approach [3] was used to validate the models derived during the staged development process. “Validation” in this context refers to the activity of gaining confidence that the formal models developed are consistent with the requirements expressed in the requirements document [8]. Test scenarios were defined to model interesting situations such as the absence of input pulses. These were run over the several models while collecting the test coverage data for each model. Tests developed for the

abstract model can be used, adapted, as regression tests in the later model development phases.

To validate the sequential model we used the re-shaped scenarios, augmenting them with new tests derived from the process of constructing and debugging of the model and its algorithmic subtleties. The validation process involves loading the chamber senses into the `Environment` which will deliver them at the correct time to the respective lead. During the simulation reactions (pulses delivered by the leads) were collected by the environment and then displayed. All of the test scenarios were reused in the validation of the *Concurrent* model and the *DR-T* model using the same paradigm.

5.2 Timing Conjectures and their Validation

The capabilities of VDMTools have been extended to support automated checking of timing-related conjectures on traces derived from runs of test scenarios over VDM++ models [6]. A simple language of standard conjecture forms has been defined and the semantics have been embedded directly into the tool set. The result of checking conjectures on a trace is displayed using the trace display format, with conjecture violation points identified as shown in Figure 2. These timing conjectures are not part of the requirements; they are assertions that the developers expect to hold over the traces derived from scenario executions.

The forms of timing conjecture relevant to the pacemaker study are: *separations*, *required separations* and *deadlines*. Separation conjectures describe a minimum separation between specified events, should the events occur. A *Separation* conjecture is a 5-tuple $separate(e_1, c, e_2, d, m)$ where e_1 and e_2 are the names of events, c is a state predicate, d is the minimum acceptable delay between an occurrence of e_1 and any following occurrence of e_2 provided that c evaluates to true at the occurrence time of e_1 . If c evaluates to false when e_1 occurs, the validation conjecture holds independently of the occurrence time of e_2 . The Boolean flag m , when set to true, indicates a requirement that the occurrence numbers of e_1 and e_2 should be equal. This allows the designer to record conjectures that describe some coordination between events. The *Required separation* conjecture is similar to the separation conjecture but additionally requires that the e_2 event does occur. The *Deadline* conjecture places a maximum delay on the occurrence of the reaction event. Again, the *match* option may be used to link the occurrence numbers of the stimulus and reaction events. A validation conjecture $deadline(e_1, c, e_2, d, m)$ consists of a stimulus event, condition and reaction event; if c holds, d is the maximum tolerable delay between stimulus and reaction.

Validation conjectures can be proposed for the test scenarios on the Distributed Real-Time model. For example, a conjecture might be stated that the minimum delay between a ventricular pace event and the next ventricular pace shall be 5000ms. This is expressed as the following conjecture C1:

```
separate(#fin(Lead`dischargePulse(-, <VENTRICLE>), true,
         #fin(Lead`dischargePulse(-, <VENTRICLE>), 5000, false)
```

A requirement that, after an atrial event there must be a ventricular pace after 1500ms (± 4 ms), leads to the following conjecture which includes a requirement that the second event occurs, C2:

```
separatereq(#fin(Lead`dischargePulse(-,<ATRIUM>), true,  
            #fin(Lead`dischargePulse(-,<VENTRICLE>), 1460, true)
```

A requirement on the maximum delay between pulses being, say, 1000 ± 8 ms would be expressed as a deadline conjecture as follows, C3:

```
deadline(#fin(Lead`dischargePulse(-,<ATRIUM>), true,  
          #fin(Lead`dischargePulse(-,<ATRIUM>), 1008, false)
```

The three validation conjectures above have been applied to test runs of the validation scenarios. In several cases this identified violations in the model and in this way the model could be improved as such bottlenecks was discovered.

6 Related Work

Efforts are being made to support the incremental development of formal models, but this approach has not so far been extended to model-oriented specifications of real-time systems with explicit deployment. Work in SCTL/MUST[22] addresses the iterative production of early-stage models of real-time systems. As in our approach, validation by testing is supported and the model production process feeds back into requirements scenarios. The Credo project [23] focuses on modelling and analysis of evolutionary structures for distributed services and also includes formal models, in a combination of Creol and Reo, similar to those described here but without so far considering deployment issues. The incremental approach suggested here also has similarities with refinement-oriented approaches, such as those in event-based B work [24] but here the focus is more on the formal aspects of the refinement does not explicitly address time or deployment.

Related work has been done by Suhaib et al. [25] in proposing a methodology derived from that of eXtreme Programming, in which “user stories” are expressed as LTL formulae representing properties which are model-checked. On each iteration, new user stories are addressed. The ordering of properties is significant for the practical tractability of the analysis on each iteration. In the context of research on real-time UML [26], a combination of UML and SDL [27] with a rigorous semantic foundation. However, in this work the ability to carry out the validation is more limited when deployment is considered. Burmester et al. [28] describe support for an iterative development process for real-time system models in extended UML by means of compositional model checking, and Uchitel et al. [29] address the incremental development of message sequence charts, again model-checking the models developed in each iteration.

7 Concluding Remarks and Further Work

Our objective in the work reported here was to assess the feasibility of using an incremental approach in the production of a useful model of a realistic real-time distributed system. The pacemaker case study suggests that such an approach can yield a viable model that can be subjected to useful validation against system-level properties at an early stage in the development process. The study encourages us to apply the approach

to a wider range of examples. Of the 19 modes of the pacemaker, eight have been modelled so far, covering 18 of the 26 controlling variables. The study revealed that the regression test suite built from the validation activities on the intermediate models was valuable in validating the later, more complex models.

Our approach has been very pragmatic, driven by the aim of providing a fully formal modelling approach with a low barrier to industrial adoption. We have not yet dealt with the relationship between the incremental addition of detail and formal refinement. In particular, we would like to be able to drive useful proof obligations out of the “refinement” steps. An examination of this issue must address the treatment of atomicity in the abstract and sequential models (for example in handling the maintenance of invariants). To encourage adoption, we feel it is essential that we automate a larger part of the validation process.

Acknowledgments We are grateful to Brian Larson for helpful suggestions on drafts of this work; to José Nuno Oliveira, Shin Sahara, Marcel Verhoef and Sander Vermolen for assisting us with the development of the extensions for the development of embedded distributed real-time systems as well as the methodology described in this paper.

References

1. Jones, C.B.: Systematic Software Development Using VDM. Second edn. Prentice-Hall International, Englewood Cliffs, New Jersey (1990) ISBN 0-13-880733-7.
2. Fitzgerald, J., Larsen, P.G.: Modelling Systems – Practical Tools and Techniques in Software Development. Cambridge University Press, The Edinburgh Building, Cambridge CB2 2RU, UK (1998) ISBN 0-521-62348-0.
3. Fitzgerald, J., Larsen, P.G., Mukherjee, P., Plat, N., Verhoef, M.: Validated Designs for Object-oriented Systems. Springer, New York (2005)
4. Fitzgerald, J., Larsen, P.G., Sahara, S.: VDMTools: advances in support for formal modeling in VDM. Sigplan Notices (Submitted for publication 2007)
5. Verhoef, M., Larsen, P.G., Hooman, J.: Modeling and Validating Distributed Embedded Real-Time Systems with VDM++. In Misra, J., Nipkow, T., Sekerinski, E., eds.: FM 2006: Formal Methods, Lecture Notes in Computer Science 4085 (2006) 147–162
6. Fitzgerald, J.S., Larsen, P.G., Tjell, S., Verhoef, M.: Validation Support for Real-Time Embedded Systems in VDM++. Technical Report CS-TR-1017, School of Computing Science, Newcastle University (April 2007) Revised version to appear in Proc. 10th IEEE High Assurance Systems Engineering Symposium, November, 2007, Dallas, Texas, IEEE.
7. CSK: Development Guidelines for Real Time Systems using VDMTools. Technical report, CSK (2006)
<http://www.vdmttools.jp/uploads/manuals/vicedevelopmentprocessE.pdf>.
8. Boston Scientific: Pacemaker system specification. Technical report, Boston Scientific (January 2007) http://www.cas.mcmaster.ca/sqrl/_SQRLDocuments/PACEMAKER.pdf.
9. Woodcock, J.: First steps in the verified software grand challenge. Computer **39**(10) (2006) 57–64
10. Fitzgerald, J., Jones, C.: Proof in the validation of a formal model of a tracking system for a nuclear plant. In Bicarregui, J., ed.: Proof in VDM: Case Studies. FACIT Series. Springer-Verlag (1998)

11. Mukherjee, P., Bousquet, F., Delabre, J., Paynter, S., Larsen, P.G.: Exploring Timing Properties Using VDM++ on an Industrial Application. In Bicarregui, J., Fitzgerald, J., eds.: Proceedings of the Second VDM Workshop. (September 2000) Available at www.vdmportal.org.
12. P. G. Larsen and B. S. Hansen and H. Brunn N. Plat and H. Toetenel and D. J. Andrews and J. Dawes and G. Parkin and others: Information technology – Programming languages, their environments and system software interfaces – Vienna Development Method – Specification Language – Part 1: Base language (December 1996)
13. Overture Group: The VDM Portal. <http://www.vdmportal.org> (2007)
14. Lano, K.: Logic specification of reactive and real-time systems. *Journal of Logic and Computation* **8**(5) (1998) 679–711
15. Verhoef, M., Larsen, P.G.: Enhancing VDM++ for Modeling Distributed Embedded Real-time Systems. Technical Report (to appear), Radboud University Nijmegen (March 2006) A preliminary version of this report is available on-line at <http://www.cs.ru.nl/~marcelv/vdm/>.
16. Verhoef, M., Larsen, P.G.: Interpreting Distributed System Architectures Using VDM++ – A Case Study. In Sauser, B., Muller, G., eds.: 5th Annual Conference on Systems Engineering Research. (March 2007) Available at <http://www.stevens.edu/engineering/cser/>.
17. CSK: VDMTools homepage. <http://www.vdmtools.jp/en/> (2007)
18. Kurita, T., Oota, T., Nakatsugawa, Y.: Formal specification of an embedded IC for cellular phones. In: Proceedings of Software Symposium 2005, Software Engineers Associates of Japan (June 2005) 73–80 (in Japanese).
19. Fitzgerald, J.S., Larsen, P.G.: Triumphs and Challenges for the Industrial Application of Model-Oriented Formal Methods. In Margaria, T., Philippou, A., Steffen, B., eds.: Proc. 2nd Intl. Symp. on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA 2007). (2007) Also Technical Report CS-TR-999, School of Computing Science, Newcastle University.
20. Macedo, H.: Validating and Understanding Boston Scientific Pacemaker Requirements. Master’s thesis, Minho University, Portugal (October 2007)
21. Macedo, H.: VDM models of the Pacemaker Challenge (2007) <http://www.vdmportal.org/twiki/bin/view/Main/PacemakerCaseStudy>.
22. Vilas, A.F., Arias, J.J.P., Redondo, R.P.D., Martinez, A.B.B.: Formalizing Incremental Design in Real-time Area: SCTL/MUS-T. In: Proceedings of the 26 th Annual International Computer Software and Applications Conference (COMPSAC02), IEEE (2002)
23. de Boer, F.: CREDO: Modeling and analysis of evolutionary structures for distributed services (2007) <http://www.cwi.nl/projects/credo/>.
24. Lecomte, T.: Event B Reference Manual. Technical report, MATISSE/ClearSy (2001)
25. Suhaib, S.M., Mathaikutty, D.A., Shukla, S.K., Berner, D.: XFM: An Incremental Methodology for Developing Formal Models. *ACM Transactions on Design Automation of Electronic Systems* **10**(4) (October 2005) 589–609
26. Douglas, B.P.: Real Time UML – Advances in the UML for real-time systems. Third edn. Addison-Wesley (2004)
27. de Jong, G.: A UML-Based Design Methodology for Real-Time and Embedded Systems. In: Proceedings of the 2002 Design, Automation and Test in Europe Conference and Exhibition (DATE.02), IEEE (2002)
28. Burmester, S., Giese, H., Hirsch, M., Schilling, D.: Incremental Design and Formal Verification with UML/RT in the FUJABA Real-Time Tool Suite. In: Proceedings of the International Workshop on Specification and validation of UML models for Real Time and embedded Systems, SVERTS2004, UML2004 (2004)
29. Uchitel, S., Kramer, J., Magee, J.: Incremental Elaboration of Scenario-Based Specifications and Behavior Models Using Implied Scenarios. *ACM Transactions on Software Engineering and Methodology* **13**(1) (January 2004) 37–85