

SOFTWARE VERIFICATION RESEARCH CENTRE

THE UNIVERSITY OF QUEENSLAND

**Queensland 4072
Australia**

TECHNICAL REPORT

No. 98-14

**Separating Timing and Calculation
in Real-Time Refinement**

Ian Hayes

July 1998

Phone: +61 7 3365 1003

Fax: +61 7 3365 1533

<http://svrc.uq.edu.au>

Copyright © 1998 Springer-Verlag. Published in Jim Grundy, Martin Schwenke and Trevor Vickers, editors, *International Refinement Workshop and Formal Methods Pacific 1998*, pages 1–16, Springer-Verlag Series in Discrete Mathematics and Theoretical Computer Science, Singapore, 1998.

Note: Most SVRC technical reports are available via anonymous ftp, from [svrc.uq.edu.au](ftp://svrc.uq.edu.au) in the directory `/pub/techreports`. Abstracts and compressed postscript files are available from <http://svrc.uq.edu.au>

Separating Timing and Calculation in Real-Time Refinement

Ian Hayes

Abstract.

We consider the specification and refinement of sequential real-time programs. Our real-time specifications describe the allowable behaviours of an implementation in terms of the values of variables over time. Hence within a specification the values of the variables and the times at which they have those values are intertwined. However, in a real-time program some commands are concerned with calculating the right outputs, while other commands, such as delays and deadlines, are concerned with making sure the outputs appear at the right time.

During the refinement process we would like to decompose the overall problem into those aspects dealing with time and those that are purely calculation. We need refinement rules that allow us to separate these concerns. Further, given a component that is only concerned with calculation, the complexities of the real-time calculus that deal with timing behaviour are an unnecessary burden. Such calculational components can be developed more straightforwardly in the standard refinement calculus. We would like to allow the use of the untimed calculus for the development of such components. To do that we need to embed the untimed calculus within the real-time calculus.

1 Introduction

We consider the problem of developing a sequential real-time program via a process of stepwise refinement from a specification of the required real-time behaviour of a system. The program operates in an environment where it can sample *inputs*, perform calculations, and update *outputs*. Of course, the all important extra dimension to the problem is time. Not only must the program calculate the correct values to be output, it must also meet the specified timing requirements.

Following the approach used in the physical sciences for many centuries, we consider inputs and outputs to be traces over time (i.e., functions from time to their value at that time). A specification determines the allowable values of the outputs in terms of the values of the inputs. However, because inputs and outputs are both traces over time, a specification inherently includes both timing and value considerations rolled into one. As part of our refinement process, we would like to decompose the problem into those aspects that deal with real-time constraints and those that perform calculations. The aim of our approach is to be able to develop the components that perform calculations using essentially the standard (untimed) sequential refinement calculus. To do this, we need to be able to perform refinement steps that separate

out timing constraints, leaving specifications of calculational components. To use standard refinement calculus derivations for the calculational components, we need to embed the standard calculus within the real-time calculus.

In Section 2 we provide an overview of the sequential real-time refinement calculus [3]. Section 3 introduces an example to illustrate the goals of the rest of the paper, and Sections 4 and 5 look at the relationship between the timed and untimed calculi, and the embedding of the untimed calculus within the real-time calculus.

2 Sequential real-time refinement

Environment In the real-time calculus, variables are modelled as functions over time. For example, the variable v of type T is modelled by a function from $Time$ to T

$$v : Time \rightarrow T$$

where $Time$ is represented by nonnegative reals. In the environment, ρ , of a program we distinguish between inputs, outputs and local variables. Inputs are under the control of the environment, not the program, although the program may make assumptions about the behaviour of inputs. Both outputs and local variables are under the control of the program; collectively we refer to outputs and local variables as the program variables and use the abbreviation $\hat{\rho}$ to stand for the program variables in an environment ρ . Outputs are distinguished from local variables because outputs are externally observable, and hence refinements must take care to preserve their behaviour.

The special variable τ is used to stand for the current time. It is not itself a function of time.

The target programming language consists of the guarded command language [1] extended with the following real-time commands:

- **delay until** D – delays execution until the absolute time D ;
- $x : \mathbf{gettime}$ – assigns the current time to x ;
- $x : \mathbf{read}(v)$ – samples the input v and assigns the value to x ; and
- **deadline** D – specifies an absolute time deadline, D , to be met by the preceding code.

Of these the deadline command [2] is novel. It cannot be implemented by simply generating machine code (generating machine code only makes it harder to meet a deadline). Instead, a timing analysis phase is required to ensure that for every valid execution of the machine code, its deadlines will always be met. If the timing analysis is successful then the machine code is a valid implementation, otherwise it must be discarded.

Wide-spectrum language As with the standard refinement calculus, our language is extended with a specification command. However, in our case the specification is dealing with variables that are traces over time as well as the special variable, τ , that stands for the current time. The form of a specification command is $\star \tilde{x} : [R]$. The frame, \tilde{x} , is the vector of variables that may be modified by the command. Any program variables in the environment that are not in

the frame remain unchanged (stable) for the duration of the execution of the command. We introduce the abbreviation

$$\text{stable}(y, S) == (\forall t1, t2 : S \bullet y(t1) = y(t2))$$

in order to state such stability properties. We also allow a vector of variables to be used in place of y , in which case all of the variables in the vector are stable over S .

The postcondition, R , of a specification command is a predicate specifying the effect that the command is to achieve. The real-time specification command has been defined by Utting and Fidge [7] by encoding it as a standard specification command that changes only time:

$$\star \tilde{x} : [R] == \tau : [R \wedge \text{'}\tau \leq \tau' \wedge \text{stable}(\hat{\rho} \setminus \tilde{x}, [\text{'}\tau \dots \tau'])] \quad (1)$$

where $\text{'}\tau$ and τ' stand for the start and finish times of the command (equivalent to Morgan's τ_0 and τ convention [6]), and $\hat{\rho} \setminus \tilde{x}$ is the set of program variables in the environment ρ minus the frame \tilde{x} . As well as establishing R , the specification also ensures that time does not go backward, and that the program variables that are not in the frame are stable for the duration of the command.

Stable predicates Assertions within real-time specifications may refer to inputs and to the current time. Because of this, one cannot guarantee that just because an assertion, P , holds at time t , that P will hold at any later time, even if the program has not modified any of the variables occurring in P in the intervening period. For example, the assertion $\tau \leq 1$ holds at time 1 but does not hold at any later time. On the other hand, if the assertion $1 \leq \tau$ holds at some time t , it holds for all later times. If v is an input then, if the assertion $v(\tau) = 0$ holds at time 1, that does not guarantee that it holds at later times, because the value of the input may change independently of the action (or lack thereof) of the program.

In dealing with real-time programs, assertions that are stable over time have special significance. For example, consider the following refinement, for introducing a selection command, that is valid in the standard refinement calculus

$$\begin{aligned} \star\{P\}; \star x : [R] &\sqsubseteq \text{if } B \rightarrow \star\{P \wedge B\}; \star x : [R] \\ &\quad \parallel \neg B \rightarrow \star\{P \wedge \neg B\}; \star x : [R] \\ &\text{fi} \end{aligned}$$

where B is some boolean expression. The above refinement is not necessarily valid in the timed calculus, because evaluating the guards takes time. Although P may be assumed to hold at the start time of the selection command, it may no longer hold at the beginning of a command within a branch because time has passed during the evaluation of the guards.

To avoid this problem, we insist that P is an *idle-stable* predicate, i.e., P does not change its value if only the current time and the value of inputs change. If we assume that, within P , $\text{'}\tau$ is used to reference the current time, then P is idle-stable provided

$$\text{'}\tau \leq \tau' \wedge \text{stable}(\hat{\rho}, [\text{'}\tau \dots \tau']) \wedge P \Rightarrow P[\text{'}\tau \setminus \tau']$$

That is, if P holds at time $\text{'}\tau$, and all the program variables are stable from $\text{'}\tau$ until some later time, τ' , then P still holds at τ' .

A further problem arises with the guards themselves. If the guards refer to inputs or to the current time, then their values are not necessarily stable. Although the guard B may evaluate to true, there is no guarantee that B will still hold at the time the corresponding branch commences execution. To avoid this problem, we insist that guards do not refer to inputs or to the current time, and hence they are idle-stable.

Similarly, just because both branches of the selection establish R , it does not follow that the whole selection will establish R . There are two problems here:

- a branch may have established R , but it takes time for the selection command to exit after completing the branch, during which time R may be invalidated;
- if R refers to initial values of variables, then within the specification command on the left side of the refinement the initial values refer to the values at the start time of the whole selection command, whereas within a branch the initial values refer to the values of the variables when the body of the branch commenced. References to either time or inputs may have changed.

To avoid these problems, we require that such predicates are both *pre-idle-stable* and *post-idle-stable*. Within an effect predicate, R , $\backslash\tau$ is used to refer to the start time of the command, and τ' is used to refer to the finish time of the command. R is pre-idle-stable provided

$$\forall u : Time \bullet u \leq \backslash\tau \leq \tau' \wedge stable(\hat{\rho}, [u \dots \backslash\tau]) \wedge R \Rightarrow R[\backslash\tau \backslash u]$$

That is, if the program variables are stable from some time u (representing the start time of the selection command in our example) until $\backslash\tau$, and R holds with respect to start time $\backslash\tau$ and finish time τ' , then R also holds with the start time replaced by u .

R is post-idle-stable provided

$$\forall u : Time \bullet \backslash\tau \leq \tau' \leq u \wedge R \wedge stable(\hat{\rho}, [\tau' \dots u]) \Rightarrow R[\tau' \backslash u]$$

That is, if R holds for start time $\backslash\tau$ and finish time τ' , and the program variables are stable from τ' until a later time u (representing the finish time of the whole selection command), then R also holds with the finish time replaced by u .

In the selection command example, in the time periods taken to evaluate the guards and exit the selection, the program variables are stable. If the effect R is both pre and post-idle-stable, then provided R is achieved for each of the branches, it will be achieved for the whole selection command.

3 An example

Consider the example of trying to determine the size of an object passing through a sensor beam on a conveyor belt moving at constant speed. (This example is taken from [4] where a more complete refinement may be found.) The sensor signal rises when the object interrupts the beam and falls when the object has passed through the beam. The size of the object is $(f - r) * s$, where r and f are the times at which the sensor rises and falls, and s is the speed of the belt.

```

con  $r, f : Time$ ;
const  $s == 10 \text{ m/s}$  – speed in metres per second

```

The logical constants, r and f , denote the exact times at which the sensor rises and falls. We cannot use these directly within the final program. Instead we need to determine approximations, rt and ft , to r and f , and then calculate the approximate size of the object in terms of rt and ft .

```

var  $size : \mathbb{N} \mu m$ ;  – size in micrometres
var  $rt, ft : \mathbb{N} \mu s$  – times in microseconds
    
```

Although these variables have been declared to be of type natural number (with units), recall that in the underlying model they are viewed as functions from *Time* to natural numbers, and hence to determine the value of a variable at a particular time, one must index the variable by that time.

Let us concentrate on the calculation component of the refinement, assuming the approximations to r and f have been determined to within an error bound of $100 \mu s$:

$$\begin{aligned}
 &\mathbf{const} \ e == 100 \mu s; \\
 &\star \{ rt(\tau) \in [r \dots r + e] \wedge ft(\tau) \in [f \dots f + e] \}
 \end{aligned} \tag{2}$$

An assertion states a condition that is true at the point at which it appears in a program. An assertion may refer to the current time, but as the assertion takes no time, both the start time, τ , and the finish time, τ' , of an assertion are the same. To allow for simpler expression of laws later, we follow the convention of only ever using τ to refer to the current time within an assertion.

The task is to calculate the size of the object to within 1 mm, and to complete the calculation within $100 \mu s$ of the object having passed:

$$\star size: [size(\tau') \in (f - r) * s \pm 1 \text{ mm} \wedge \tau' \leq f + 100 \mu s] \tag{3}$$

To satisfy this specification, one must both calculate the size and meet the deadline. Here we can make use of the deadline command to separate calculation and timing requirements. This is a crucial step in the process as it allows one to concentrate on the calculation of the size without having to worry explicitly about the timing constraint. We make use of the *separate deadline* law [3].

Law 1 (separate deadline) *Given an environment, ρ , provided D is a time-valued expression, which may only include references to logical constants and program variables but no references to τ*

$$\star x: [P, \quad R \wedge \tau' \leq (D @ \tau')] \sqsubseteq \star x: [P, \quad R]; \mathbf{deadline} \ D$$

The expression D may contain references to program variables. These are to be interpreted as the value of the variables at the finish time of the command. The notation $D @ \tau'$ stands for the expression D with any reference to a program variable, v , replaced by $v(\tau')$. Refining (3) using this law gives

$$\begin{aligned}
 &\star size: [size(\tau') \in (f - r) * s \pm 1 \text{ mm}]; \\
 &\mathbf{deadline} \ f + 100 \mu s
 \end{aligned} \tag{4}$$

The specification command can now be refined to an assignment command using the assignment introduction rule of the real-time calculus.

Law 2 (assignment) *Given an environment, ρ , a frame, \tilde{x} , such that $\tilde{x} \subseteq \hat{\rho}$, and a vector of idle-stable expressions, \tilde{E} , provided*

$$\backslash\tau \leq \tau' \wedge \text{stable}(\hat{\rho} \setminus \tilde{x}, [\backslash\tau \dots \tau']) \wedge P \wedge \tilde{x}(\tau') = (\tilde{E} @ \backslash\tau) \Rightarrow R$$

for all states, then

$$\star\{P\}; \star\tilde{x}: [R] \sqsubseteq \tilde{x} := \tilde{E}.$$

The specification command (4) can be refined using this law to the assignment

$$\text{size} := (ft - rt) * s$$

with the following proof obligation

$$\left(\begin{array}{l} \backslash\tau \leq \tau' \wedge \\ \text{stable}(rt, [\backslash\tau \dots \tau']) \wedge \\ \text{stable}(ft, [\backslash\tau \dots \tau']) \wedge \\ rt(\backslash\tau) \in [r \dots r + e] \wedge \\ ft(\backslash\tau) \in [f \dots f + e] \wedge \\ \text{size}(\tau') = (ft(\backslash\tau) - rt(\backslash\tau)) * s \end{array} \right) \Rightarrow \text{size}(\tau') \in (f - r) * s \pm 1 \text{ mm} \quad (5)$$

If we consider the refinement of the specification command (4) in the standard untimed refinement calculus, the assumptions (2) can be written

$$\star\{\backslash rt \in [r \dots r + e] \wedge \backslash ft \in [f \dots f + e]\} \quad (6)$$

where $\backslash rt$ stands for the value of the variable at the time at which the assertion is reached. We use the before-state decoration to simplify writing laws involving assertions as preconditions. The specification (4) can be written as

$$\star\text{size}: [\text{size}' \in (f - r) * s \pm 1 \text{ mm}] \quad (7)$$

where size' refers to the value of the variable size on termination of the command. This can be refined to the same assignment as before but with the following proof obligation:

$$\left(\begin{array}{l} \backslash rt \in [r \dots r + e] \wedge \\ \backslash ft \in [f \dots f + e] \end{array} \right) \Rightarrow (\backslash ft - \backslash rt) * s \in (f - r) * s \pm 1 \text{ mm} \quad (8)$$

In this case the standard refinement contains all the essential information of the refinement in the timed calculus, but omits the complexities of explicit time indices and stability conditions. More importantly the proof obligation is considerably simpler. The goal of this paper is to show how this simpler approach for calculational components can be embedded within the real-time refinement calculus.

Notational conventions We start by introducing a notational convention that allows us to write the simpler form of specification within the real-time calculus. Within assertions and specifications, we allow a variable, v , to be referenced in one of three ways:

- v refers to the whole trace of v (it may be explicitly indexed to get the value of v at a particular time);
- $\backslash v$ refers to the value of v at the current time within an assertion, and within a specification to the value of v at the start time of the specification; and
- v' , which is only meaningful in a specification, refers to the value of v at the finish of the command.

Using this convention within assertions and specifications, $\backslash v$ is equivalent to $v(\backslash \tau)$, and within a specification v' is equivalent to $v(\tau')$. Hence the assertions (2) and (6) are equivalent, and the specifications (4) and (7) are equivalent.

Our choice of notation is motivated by wanting v' and $\backslash v$ to have the same meaning as v and v_0 in Morgan's refinement calculus [6], while allowing references to the value of v at times other than the start and finish times of a specification using the trace notation v . Our choice of notation for this paper is based on the desire to have an identifier v stand for its whole timed trace, rather than its value at some particular time. Then to avoid ambiguities, we need separate notations for the before and after-state values of variables; the notation chosen is derived from that of Hehner [5]. (Aside: In earlier papers we have used the notation v_0 to stand for the initial value of v , and v to stand for both the whole trace and the final value of v . Unfortunately, expressions like $x = y$, where x and y are variables, are then ambiguous: the x and y could either both be whole traces or final state values, and the context cannot disambiguate the two interpretations.)

For example, the read command samples its input at some time during its execution. It can be defined in terms of a specification command as follows:

$$x : \mathbf{read}(v) == \star x: [x \in v \mid [\backslash \tau \dots \tau'] \mid]$$

where $v \mid [\backslash \tau \dots \tau'] \mid$ is the set of values of v over the closed interval from the start of the command, $\backslash \tau$, to the finish of the command, τ' . Using just the $\backslash v$ and v' conventions one cannot specify the read command because one can only refer to the values of v at $\backslash \tau$ and τ' , and not over the range in between.

Our notational conventions allow us to abbreviate assertions and specifications as illustrated above. However, they do not justify the abbreviation of the proof obligation (5) to (8). If we expand (8) using our notational conventions we get

$$\left(\begin{array}{l} rt(\backslash \tau) \in [r \dots r + e] \wedge \\ ft(\backslash \tau) \in [f \dots f + e] \end{array} \right) \Rightarrow (ft(\backslash \tau) - rt(\backslash \tau)) * s \in (f - r) * s \pm 1 \text{ mm} \quad (9)$$

which is simpler than (5). For such calculational components, the additional premisses in (5) are not needed to prove the obligation.

4 Relating the timed and untimed calculi

In order to support the use of the simpler untimed proof obligations for refinement steps for calculational components, we need to examine the relationship between the untimed and timed refinement calculi. In this section, we examine an example refinement law and its proof obligation, and show how the proof obligation for the real-time law can be transformed into a proof obligation in the untimed model.

Consider the following refinement in the real-time calculus

$$\star\{P\}; \star\tilde{x}: [Q] \sqsubseteq \star\{P\}; \star\tilde{x}: [R] \quad (10)$$

where P , Q and R are predicates in which there is no use of the $\backslash v$ and v' notation, except for $\backslash \tau$ and τ' , i.e., all references to a variable, v , treat it explicitly as a function of time. In addition, P does not refer to τ' , only $\backslash \tau$. The refinement (10) is valid provided that

$$\backslash \tau \leq \tau' \wedge \text{stable}(\hat{\rho} \backslash \tilde{x}, [\backslash \tau \dots \tau']) \wedge P \wedge R \Rightarrow Q \quad (11)$$

holds for all timed states. This is essentially just the strengthen postcondition rule of Morgan [6] adapted to the slightly different notation, and with the additional premisses $\backslash \tau \leq \tau'$ and $\text{stable}(\hat{\rho} \backslash \tilde{x}, [\backslash \tau \dots \tau'])$ derived from the definition of the specification command (1).

We would like to make use of the $\backslash v$ and v' notation to stand for $v(\backslash \tau)$ and $v(\tau')$, respectively. We can always eliminate the abbreviations by replacing $\backslash v$ by $v(\backslash \tau)$ and v' by $v(\tau')$. We define the notation $R @ (\backslash \tau, \tau')$ to be R with all references to $\backslash v$ replaced by $v(\backslash \tau)$ and all references to v' replaced by $v(\tau')$, for all variables, v , occurring in R .

If the predicates P , Q and R in the above refinement rule are allowed to use the abbreviated notation (noting that P is restricted to use only $\backslash v$ and not v'), then (11) has to be rewritten

$$\backslash \tau \leq \tau' \wedge \text{stable}(\hat{\rho} \backslash \tilde{x}, [\backslash \tau \dots \tau']) \Rightarrow (P \wedge R \Rightarrow Q) @ (\backslash \tau, \tau') \quad (12)$$

Now consider the special case of a calculational component in which P , Q and R make no explicit references to the trace variable v , i.e., they only reference v via the $\backslash v$ and v' conventions. Further, P , Q and R make no references to $\backslash \tau$ and τ' . For this case P is essentially a predicate constraining the pre-state values of variables, and Q and R are predicates relating the pre and post-state values of variables.

The stability predicate implies that for every program variable, v , that is not in the frame ($v \in \hat{\rho} \backslash \tilde{x}$), its value does not change and therefore we may deduce that $\backslash v = v'$. This allows one to replace any occurrences of v' with $\backslash v$, and hence to rewrite (12) as

$$\left(\begin{array}{l} \backslash \tau \leq \tau' \wedge \\ \text{stable}(\hat{\rho} \backslash \tilde{x}, [\backslash \tau \dots \tau']) \end{array} \right) \Rightarrow (P \wedge R \Rightarrow Q) [v' \backslash \backslash v]_{v \in \hat{\rho} \backslash \tilde{x}} @ (\backslash \tau, \tau') \quad (13)$$

where the notation $[v' \backslash \backslash v]_{v \in \hat{\rho} \backslash \tilde{x}}$ stands for the replacement of all occurrences of v' by $\backslash v$, for all identifiers, v , in the set $\hat{\rho} \backslash \tilde{x}$. Now we note that the time progress and stability premisses (to the left of the first implication) in (13) are of no help in establishing the right side of (13) because the predicate makes no references to $\backslash \tau$ and τ' , or to the values of variables other than via $\backslash v$ and v' . Hence we can simplify (13) to

$$(P \wedge R \Rightarrow Q) [v' \backslash \backslash v]_{v \in \hat{\rho} \backslash \tilde{x}} @ (\backslash \tau, \tau') \quad (14)$$

If (14) holds for all values of the timed traces of its variables, then the refinement (10) given above is valid. But note that the predicate to the left of the '@' is essentially that used in the untimed refinement calculus for the strengthen postcondition law (aside from the different notational conventions). If the predicate to the left of the '@' holds for all values of $\backslash v$ and v' , then (14) will hold for all timed traces. Hence all we need to show is that for all values of the before and after states

$$(P \wedge R \Rightarrow Q) [v' \backslash \backslash v]_{v \in \hat{\rho} \backslash \tilde{x}} \quad (15)$$

In the next section, we formalise the above reasoning by showing that for any predicate M that does not refer to $\backslash\tau$ or τ' and only references variables via the $\backslash v$ and v' convention, that if M holds for all untimed states, then $M @ (\backslash\tau, \tau')$ holds for all timed trace states.

5 Embedding the untimed calculus

In this section we formalise the relationship between the timed and untimed models. In the timed model, variables are represented by traces over time, while in the untimed model only the before and after-state values of variables are available. The before-state value of a variable, $\backslash v$, is the value of v at time $\backslash\tau$ and the after-state value, v' , is the value of v at time τ' . Below we formalise timed and untimed states and the relationship between them, and then we show the relationship between untimed and timed proof obligations.

In a specification command, the effect predicate may refer to logical constants, program constants, variables (inputs, outputs, and local variables), and the start and finish time of the command ($\backslash\tau$ and τ'). In both the timed and untimed models, constants are treated in the same manner, and hence we shall ignore them in this exposition. We represent environments by the following schema

<i>Environment</i>	
$in, out, lvar, pvar, vars : \mathbb{P} \textit{Ident}$	
$typeof : \textit{Ident} \rightarrow \mathbb{P} \textit{Val}$	
$pairwise_disjoint\langle in, out, lvar \rangle$	
$pvar = out \cup lvar$	
$vars = in \cup out \cup lvar \wedge \backslash\tau' \notin vars$	
$\text{dom } typeof = vars$	

where *Ident* is the set of all identifiers (including $\backslash\tau$) and *Val* is the set of all values (including booleans, integers, *Time*, etc.). The disjoint sets *in*, *out*, and *lvar* represent the names of the inputs, outputs and local variables, respectively. The set *pvar* gives the names of the program variables; for an environment, $\rho, \hat{\rho}$ is an abbreviation for $\rho.pvar$. The set *vars* gives the names of all of the variables, and the function *typeof* gives their types.

In the timed model, variables are represented by functions of time, and the start and finish times are just time values. We can model a *timed state* by a mapping from identifiers to values, in which the identifiers $\backslash\tau$ and τ' are elements of the domain of a timed state and their values are of type *time*, and variables are modelled by timed traces. Given an environment, ρ , the set of all timed states is given by

$$TState_\rho = \{ \gamma : \textit{Ident} \rightarrow \textit{Value} \mid \\ \{ \backslash\tau, \tau' \} \subseteq \text{dom } \gamma \wedge \{ \gamma(\backslash\tau), \gamma(\tau') \} \subseteq \textit{Time} \wedge \\ \rho.vars \subseteq \text{dom } \gamma \wedge (\forall id : \rho.vars \bullet \gamma(id) \in (\textit{Time} \rightarrow \rho.typeof(id))) \}$$

where *Value* includes *Val* as well as timed traces of type $\textit{Time} \rightarrow \textit{Val}$. For example, the following timed state represents a start time of zero and a finish time of two, and a variable x that has value zero up to time one, and then has value one.

$$\gamma_1 = \{ \backslash\tau \mapsto 0, \tau' \mapsto 2, x \mapsto (\lambda t : \textit{Time} \bullet \text{if } t < 1 \text{ then } 0 \text{ else } 1) \}$$

In the untimed model, we only have to model the values of variables in the before and after states. Given a plain identifier, id , we use $\backslash id$ to refer before-state decorated identifier and id' to refer to the after-state decorated identifier. The sets of all before and after-state decorated identifiers are given by $\backslash Ident$ and $Ident'$:

$$\begin{aligned}\backslash Ident &== \{id : Ident \bullet \backslash id\} \\ Ident' &== \{id : Ident \bullet id'\}\end{aligned}$$

The set of all *untimed states* is given by

$$\begin{aligned}UState_\rho &== \{\sigma : (\backslash Ident \cup Ident') \rightarrow Val \mid \\ &\quad \{\backslash \tau, \tau'\} \subseteq \text{dom } \sigma \wedge \{\sigma(\backslash \tau), \sigma(\tau')\} \subseteq Time \wedge \\ &\quad (\forall id : \rho.vars \bullet \{\backslash id, id'\} \subseteq \text{dom } \sigma \wedge \{\sigma(\backslash id), \sigma(id')\} \subseteq \rho.typeof(id))\}\end{aligned}$$

For example, the following represents an untimed state with a start time of zero and a finish time of two, and a variable, x , with an initial value of zero and a final value of one.

$$\sigma_1 = \{\backslash \tau \mapsto 0, \tau' \mapsto 2, \backslash x \mapsto 0, x' \mapsto 1\}$$

Given a timed state, γ , one can extract the corresponding untimed state: the before and after times, $\backslash \tau$ and τ' are the same in both models, and the before and after values of each variable in the untimed model are just the values of the variables in the timed model at the before and after times.

$$\left| \begin{array}{l} extract == TState_\rho \rightarrow UState_\rho \\ \hline extract(\gamma) = \{\backslash \tau \mapsto \gamma(\backslash \tau), \tau' \mapsto \gamma(\tau')\} \cup \\ \quad \{id : \rho.vars \bullet \backslash id \mapsto \gamma(id)(\gamma(\backslash \tau))\} \cup \\ \quad \{id : \rho.vars \bullet id' \mapsto \gamma(id)(\gamma(\tau'))\} \end{array} \right|$$

For example, if we apply *extract* to the example timed state γ_1 given above, then the result is equal to the example untimed state σ_1 .

Given a predicate P that only references variables via the $\backslash v$ and v' conventions we would like to show that if P holds for all states in an untimed interpretation, then $P @ (\backslash \tau, \tau')$ holds for all states in the timed interpretation. To formalise this, we need to give both the untimed and timed semantics of predicates. We begin by giving part of the syntax for terms (including predicates).

$$\begin{aligned}Term &::= constant \mid \backslash \tau \mid \tau' \mid \backslash id \mid id' \mid id \mid Term_0 @ (Term_1, Term_2) \mid \\ &\quad \neg Term \mid Term_0 \wedge Term_1 \mid Term_0 = Term_1 \mid \dots\end{aligned}$$

where *constant* stands for a constant, *id* is an identifier, and $Term_0$, $Term_1$ and $Term_2$ are terms.

Table 1 gives the untimed semantics for these constructs with the exception of variable traces and the '@' notation, neither of which make sense in the untimed interpretation. The untimed semantic function, M_U , given a term T and an untimed state σ , gives the value of the term in that state. For example,

Table 1. Untimed semantics of terms

$M_U : Term \mapsto (UState_\rho \rightarrow Val)$
$\forall \sigma : UState_\rho \bullet$
$M_U(constant)(\sigma) = constant$
$M_U(\backslash \tau)(\sigma) = \sigma(\backslash \tau')$
$M_U(\tau')(\sigma) = \sigma(\tau')$
$M_U(\backslash id)(\sigma) = \sigma(\backslash id)$
$M_U(id')(\sigma) = \sigma(id')$
$M_U(\neg T)(\sigma) = \neg M_U(T)(\sigma)$
$M_U(T_0 \wedge T_1)(\sigma) = (M_U(T_0)(\sigma) \wedge M_U(T_1)(\sigma))$
$M_U(T_0 = T_1)(\sigma) = (M_U(T_0)(\sigma) = M_U(T_1)(\sigma))$

where T_0 and T_1 are terms.

$$\begin{aligned}
 & M_U(\backslash \tau = \tau' \wedge \neg(\backslash x = x'))(\sigma) \\
 &= M_U(\backslash \tau = \tau')(\sigma) \wedge M_U(\neg(\backslash x = x'))(\sigma) \\
 &= (M_U(\backslash \tau)(\sigma) = M_U(\tau')(\sigma)) \wedge \neg M_U(\backslash x = x')(\sigma) \\
 &= \sigma(\backslash \tau) = \sigma(\tau') \wedge \neg (M_U(\backslash x)(\sigma) = M_U(x')(\sigma)) \\
 &= \sigma(\backslash \tau) = \sigma(\tau') \wedge \neg (\sigma(\backslash x) = \sigma(x'))
 \end{aligned}$$

Table 2 gives the semantics of terms in the timed model. The semantic function, M_T , given a term and a timed state, gives the value of the term in that state. In this case all constructs can be given a meaning. An (undecorated) identifier, id , stands for its whole timed trace value, $\gamma(id)$, and before and after decorated identifiers stand for the value of the trace, $\gamma(id)$, at the start and finish times, $\gamma(\backslash \tau)$ and $\gamma(\tau')$.

Table 2. Timed semantics of terms

$M_T : Term \mapsto (TState_\rho \rightarrow Value)$
$\forall \gamma : TState_\rho \bullet$
$M_T(constant)(\gamma) = constant$
$M_T(\backslash \tau)(\gamma) = \gamma(\backslash \tau')$
$M_T(\tau')(\gamma) = \gamma(\tau')$
$M_T(id)(\gamma) = \gamma(id)$
$M_T(\backslash id)(\gamma) = \gamma(id)(\gamma(\backslash \tau'))$
$M_T(id')(\gamma) = \gamma(id)(\gamma(\tau'))$
$M_T(T_0 @ (T_1, T_2))(\gamma) = M_{TT}(T_0)(\gamma)(M_T(T_1)(\gamma), M_T(T_2)(\gamma))$
$M_T(\neg T)(\gamma) = \neg M_T(T)(\gamma)$
$M_T(T_0 \wedge T_1)(\gamma) = (M_T(T_0)(\gamma) \wedge M_T(T_1)(\gamma))$
$M_T(T_0 = T_1)(\gamma) = (M_T(T_0)(\gamma) = M_T(T_1)(\gamma))$

where T_0 , T_1 and T_2 are terms.

In the definition of the ‘@’ notation, we allow terms to be used to specify the start and finish times. Hence, within the term to the left of the ‘@’, the start and finish times are interpreted as the values of these two terms. To define this, we introduce in Table 3 the auxiliary semantic function M_{TT} which takes an additional pair of parameters specifying the start and finish times. (The definition of M_{TT} for the ‘@’ operator is given here for completeness, but it corresponds to the case where an ‘@’ operator is used within the left operand of another ‘@’ operator and is not used in the remainder of this paper.)

Table 3. Auxiliary function for timed terms

$M_{TT} : Term \mapsto (TState_\rho \rightarrow (Time \times Time \rightarrow Value))$
$\forall \gamma : TState_\rho; t_0, t_1 : Time \bullet$
$M_{TT}(constant)(\gamma)(t_0, t_1) = constant$
$M_{TT}(\backslash \tau)(\gamma)(t_0, t_1) = t_0$
$M_{TT}(\tau')(\gamma)(t_0, t_1) = t_1$
$M_{TT}(id)(\gamma)(t_0, t_1) = \gamma(id)$
$M_{TT}(\backslash id)(\gamma)(t_0, t_1) = \gamma(id)(t_0)$
$M_{TT}(id')(\gamma)(t_0, t_1) = \gamma(id)(t_1)$
$M_{TT}(T_0 @ (T_1, T_2))(\gamma)(t_0, t_1) =$ $M_{TT}(T_0)(\gamma)(M_{TT}(T_1)(\gamma)(t_0, t_1), M_{TT}(T_2)(\gamma)(t_0, t_1))$
$M_{TT}(\neg T)(\gamma)(t_0, t_1) = \neg M_{TT}(T)(\gamma)(t_0, t_1)$
$M_{TT}(T_0 \wedge T_1)(\gamma)(t_0, t_1) = (M_{TT}(T_0)(\gamma)(t_0, t_1) \wedge M_{TT}(T_1)(\gamma)(t_0, t_1))$
$M_{TT}(T_0 = T_1)(\gamma)(t_0, t_1) = (M_{TT}(T_0)(\gamma)(t_0, t_1) = M_{TT}(T_1)(\gamma)(t_0, t_1))$

where T_0 , T_1 and T_2 are terms.

As an example of the application of the timed semantics, we consider the term used earlier, $\backslash \tau = \tau' \wedge \neg (\backslash x = x')$, but with an application of an ‘@’ operator added.

$$\begin{aligned}
& M_T((\backslash \tau = \tau' \wedge \neg (\backslash x = x')) @ (\backslash \tau, \tau'))(\gamma) \\
&= M_{TT}(\backslash \tau = \tau' \wedge \neg (\backslash x = x'))(\gamma)(M_T(\backslash \tau)(\gamma), M_T(\tau')(\gamma)) \\
&= M_{TT}(\backslash \tau = \tau' \wedge \neg (\backslash x = x'))(\gamma)(\gamma(\backslash \tau), \gamma(\tau')) \\
&= M_{TT}(\backslash \tau = \tau')(\gamma)(\gamma(\backslash \tau), \gamma(\tau')) \wedge M_{TT}(\neg (\backslash x = x'))(\gamma)(\gamma(\backslash \tau), \gamma(\tau')) \\
&= (M_{TT}(\backslash \tau)(\gamma)(\gamma(\backslash \tau), \gamma(\tau')) = M_{TT}(\tau')(\gamma)(\gamma(\backslash \tau), \gamma(\tau')) \wedge \\
&\quad \neg M_{TT}(\backslash x = x')(\gamma)(\gamma(\backslash \tau), \gamma(\tau'))) \\
&= (\gamma(\backslash \tau) = \gamma(\tau')) \wedge \neg (M_{TT}(\backslash x)(\gamma)(\gamma(\backslash \tau), \gamma(\tau')) = M_{TT}(x')(\gamma)(\gamma(\backslash \tau), \gamma(\tau'))) \\
&= (\gamma(\backslash \tau) = \gamma(\tau')) \wedge \neg (\gamma(x)(\gamma(\backslash \tau)) = \gamma(x')(\gamma(\tau'))) \\
&= false
\end{aligned}$$

Note that we can deduce in the timed model that this term is false, independently of the state γ , because if the start and finish times are the same then the values of a variable at the start and finish times must be the same. This was not possible for the same example in the untimed model for an arbitrary state, σ . However, if we assume that the untimed state, σ , in the untimed interpretation is the result of extracting an untimed state from some timed state,

γ , then we can also deduce the term is false. Let $\sigma = \text{extract}(\gamma)$, then

$$\begin{aligned} & \sigma(' \tau ') = \sigma(' \tau ') \wedge \neg (\sigma(' x ') = \sigma(' x ')) \\ & = \text{extract}(\gamma)(' \tau ') = \text{extract}(\gamma)(' \tau ') \wedge \neg (\text{extract}(\gamma)(' x ') = \text{extract}(\gamma)(' x ')) \\ & = \gamma(' \tau ') = \gamma(' \tau ') \wedge \neg (\gamma(' x ') (\gamma(' \tau ')) = \gamma(' x ') (\gamma(' \tau '))) \\ & = \text{false} \end{aligned}$$

Note that the second last line is the same as that in the timed semantics case. This example is a special case of the general theorem that we would like to prove.

Theorem 1. *For any predicate P that does not use the '@' operator, and only references variables via the $\backslash v$ and v' conventions, if P holds for all states in an untimed interpretation, then $P @ (\backslash \tau, \tau')$ holds for all states in the timed interpretation:*

$$(\forall \sigma : U\text{State}_\rho \bullet M_U(P)(\sigma)) \Rightarrow (\forall \gamma : T\text{State}_\rho \bullet M_T(P @ (\backslash \tau, \tau'))(\gamma))$$

Proof. The proof relies on Lemma 1 below.

$$\begin{aligned} & \forall \gamma : T\text{State}_\rho \bullet M_T(P @ (\backslash \tau, \tau'))(\gamma) \\ & \equiv \text{Lemma 1} \\ & \forall \gamma : T\text{State}_\rho \bullet M_U(P)(\text{extract}(\gamma)) \\ & \Leftarrow \text{change of variable; } \text{range}(\text{extract}) \subseteq U\text{State}_\rho \\ & \forall \sigma : U\text{State}_\rho \bullet M_U(P)(\sigma) \end{aligned}$$

Lemma 1. *Given any term, P , satisfying the same restrictions as for Theorem 1, the timed meaning of a term $P @ (\backslash \tau, \tau')$ in a timed state γ , is equal to the untimed meaning of P in the untimed state extracted from γ .*

$$\forall \gamma : T\text{State}_\rho \bullet M_T(P @ (\backslash \tau, \tau'))(\gamma) = M_U(P)(\text{extract}(\gamma))$$

Proof. The proof makes use of the semantics of terms in both the timed and untimed interpretations, and uses structural induction over terms.

$$\begin{aligned} & M_T(P @ (\backslash \tau, \tau'))(\gamma) \\ & = \text{timed semantics of '@'} \\ & M_{TT}(P)(\gamma)(M_T(\backslash \tau)(\gamma), M_T(\tau')(\gamma)) \\ & = \text{timed semantics of } \backslash \tau \text{ and } \tau' \\ & M_{TT}(P)(\gamma)(\gamma(' \tau '), \gamma(' \tau ')) \\ & = \text{structural induction} \\ & M_U(P)(\text{extract}(\gamma)) \end{aligned}$$

The last step is proved by structural induction over terms. Selected base cases follow:
Constants:

$$\begin{aligned} & M_{TT}(\text{constant})(\gamma)(\gamma(' \tau '), \gamma(' \tau ')) \\ & = \text{timed semantics} \\ & \text{constant} \\ & = \text{untimed semantics} \\ & M_U(\text{constant})(\text{extract}(\gamma)) \end{aligned}$$

Start time:

$$\begin{aligned}
& M_{TT}(\backslash\tau)(\gamma)(\gamma(\backslash\tau'), \gamma(\tau')) \\
& = \text{timed semantics} \\
& \quad \gamma(\backslash\tau') \\
& = \text{definition of extract} \\
& \quad \text{extract}(\gamma)(\backslash\tau') \\
& = \text{untimed semantics} \\
& \quad M_U(\backslash\tau)(\text{extract}(\gamma))
\end{aligned}$$

Before variable value:

$$\begin{aligned}
& M_{TT}(\backslash id)(\gamma)(\gamma(\backslash\tau'), \gamma(\tau')) \\
& = \text{timed semantics} \\
& \quad \gamma(id)(\gamma(\backslash\tau')) \\
& = \text{definition of extract} \\
& \quad \text{extract}(\gamma)(\backslash id) \\
& = \text{untimed semantics} \\
& \quad M_U(\backslash id)(\text{extract}(\gamma))
\end{aligned}$$

The inductive cases are straightforward. For example:

$$\begin{aligned}
& M_{TT}(T_0 \wedge T_1)(\gamma)(\gamma(\backslash\tau'), \gamma(\tau')) \\
& = \text{timed semantics} \\
& \quad M_{TT}(T_0)(\gamma)(\gamma(\backslash\tau'), \gamma(\tau')) \wedge M_{TT}(T_1)(\gamma)(\gamma(\backslash\tau'), \gamma(\tau')) \\
& = \text{inductive hypothesis} \\
& \quad M_U(T_0)(\text{extract}(\gamma)) \wedge M_U(T_1)(\text{extract}(\gamma)) \\
& = \text{untimed semantics} \\
& \quad M_U(T_0 \wedge T_1)(\text{extract}(\gamma))
\end{aligned}$$

6 Conclusions

The sequential real-time refinement calculus [2] introduced a novel deadline command as a way of including timing deadlines within machine-independent, higher-level, real-time programs. In addition, it introduced laws, such as *separate deadline*, that allow the timing related parts of a specification to be separated from the calculational parts. However refinements of the calculational components in the raw real-time calculus are encumbered by the additional baggage required for the real-time components. The objective of this paper has been to show that the untimed calculus can be embedded within the real-time calculus via the use of some special notational conventions. This allows components of programs that are only responsible for performing calculations rather than meeting real-time constraints, to be developed using the simpler untimed calculus.

The notational conventions allow abbreviated reference to the values of variables at the start and finish times of a specification command, as well as reference to the value of a variable at other times via the use of the trace notation. If a specification only uses the before and after state conventions, i.e., it is a calculational component, then the simpler untimed refinement

laws may be used. To justify this approach, we have presented both a timed and untimed semantics for predicates using this convention, and shown that if a predicate P holds for all untimed states, then the corresponding predicate in the timed model, $P@(\tau, \tau')$, holds for all states in the timed model.

The overall objective of our work is to provide a systematic method for the development of real-time programs, and to provide tool support for such a method. This paper is a step towards a method that supports both the sequential real-time refinement calculus, and embedding within that the standard sequential calculus. The use of the embedded untimed calculus allows simpler developments of program fragments that are not concerned with timing constraints.

Acknowledgements I would like to thank Ray Nickson and Peter Robinson for pointing out possible ambiguities in earlier versions of the notation, and David Carrington, Andy Evans and Luke Wildman for feedback on earlier drafts of this paper.

References

1. E. W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
2. I. J. Hayes and M. Utting. Coercing real-time refinement: A transmitter. In D. J. Duke and A. S. Evans, editors, *BCS-FACS Northern Formal Methods Workshop (NFMW'96)*. Springer, 1997.
3. I. J. Hayes and M. Utting. A sequential real-time refinement calculus. Technical Report UQ-SVRC-97-33, Software Verification Research Centre, The University of Queensland, URL <http://svrc.uq.edu.au>, 1997.
4. I. J. Hayes and M. Utting. Deadlines are termination. In D. Gries and W.-P. de Roever, editors, *IFIP TC2/WG2.2, 2.3 International Conference on Programming Concepts and Methods (PRO-COMET'98)*, pages 186–204. Chapman and Hall, 1998.
5. Eric C. R. Hehner. Predicative programming: Parts i and ii. *Comm. ACM*, 27(2):134–151, February 1984. Corrigendum: *Comm. ACM* Vol. 27 No. 6 (June 1984) p. 593.
6. C. C. Morgan. *Programming from Specifications*. Prentice Hall, second edition, 1994.
7. M. Utting and C. J. Fidge. A real-time refinement calculus that changes only time. In He Jifeng, editor, *Proc. 7th BCS/FACS Refinement Workshop*, Electronic Workshops in Computing. Springer, July 1996.