Reactive System Specification and Refinement K. Lano

Dept. of Computing, Imperial College, 180 Queens Gate, London SW7 2BZ.

Abstract. This paper describes formal approaches for reactive and real time system specification and development, using a process of systematic translation from statechart descriptions of a system into a specification language utilising real time logic (RTL), and refinement within this language. Alternative implementation strategies using synchronisation constraints and synchronisation code are also provided, together with examples of development using the approach.

The approach provides a unitary formalism which combines statecharts, RTL and temporal logic. Animation and proof tools are also briefly described.

1 Introduction

Software applications whose main purpose is to interact in real time with external systems or devices are termed 'reactive systems'. Examples include chemical process control systems, or patient monitoring systems, in which response to external events is a major part of the system requirements, and the timing of these responses is also critical. More generally, many safety critical systems are reactive or real time systems, for which formal specifications are becoming more frequently required by standards. Thus a formal language capable of treating issues of process scheduling, timing, concurrent execution and interrupts, and of supporting reasoning about such aspects at a suitable level of abstraction, is highly desirable.

In this context the use of an object-oriented paradigm is appropriate: objects may correspond to reactive components whose behaviour can be specified in isolation, or with real world concepts or entities, such as particular controlled devices, thus enabling a systematic tracing of requirements through specifications to code to be performed, supporting validation, assessment and maintenance.

The formal specification language Z^{++} [10, 6, 8] will be considered as the basis for a real time specification language. This paper extends the Z^{++} language by allowing the use of RTL formulae in the HISTORY clause of a class. In contrast to VDM⁺⁺ [10], this formalism is more orientated towards *discrete event systems* rather than continuous variable dynamic systems [14]. However, the formalism presented here could be used as a unifying semantics for much of VDM⁺⁺, and was initially defined to serve this purpose.

Section 2 gives the syntax of Z^{++} specifications, together with examples. Section 3 summarises the temporal logic formalism used in the paper. Section 4 describes the development process intended to be used with the language, and summarises the steps in this process. Section 5 introduces a refinement technique using semaphores to implement synchronisation constraints, and Section 6 describes refinement using synchronisation code [13].

2 Z⁺⁺/RTL Syntax

2.1 Classes

A Z^{++} specification consists of a collection of *class* definitions, including generic classes. A schematic Z^{++} class is of the form:

```
CLASS C[TypeParameters]
EXTENDS Ancestors
TYPES
  Tdefs
FUNCTIONS
  Axdefs
OWNS
  С
INVARIANT
  Invc
OPERATIONS
  [*] m : IN \rightarrow OUT;
   . . . .
RETURNS
  \mathbf{r} : IN \rightarrow OUT;
   . . . .
ACTIONS
   Prem.C &
          [*] m x y ==>
                                Def_{m,C};
   . . . .
HISTORY
          H_{C}
END CLASS
```

Each of the clauses in the body of a class definition are optional, although an ACTIONS clause requires an OPERATIONS clause. The **TypeParameters** are a list (possibly empty) of *generic* type parameters used in the class definition. The EXTENDS list is the set of previously defined classes that are inherited by this class. Local types, functions and constants can be defined in the TYPES and FUNCTIONS clauses. The OWNS list gives attribute declarations. The INVARIANT specifies time-invariant properties of the internal state. The default invariant is true.

The ACTIONS list gives the definitions of the various operations that can be performed on instances of the object. The default action for a method, if no action for it is listed, is the completely non-deterministic operation on the state of the class and its parameter types. Input parameters are listed before the output parameters in the action definitions. Z predicates, method invocations and the B0 procedural code constructs of Abrial's B Notation [11] can be used to define methods. Operations are given explicit preconditions by the notations

The default precondition is true. Methods with a preceeding * are *internal* actions and are discussed further below. The HISTORY of a class is an RTL predicate, the forms of which are given in Section 3.

Details of the reference semantics of Z^{++} are given in [6]. Essentially a class name C, when used as a type, denotes a countably infinite set @C of references to objects of C. A function

 $*_{\mathbf{C}} : @\mathbf{C} \rightarrow \mathbf{State}_{\mathbf{C}}$

then obtains actual object values (elements of the implicit state schema of C) from these references. \overline{C} denotes dom $*_{C}$, the set of existing objects of C.

An operation New_C creates a new instance of C and modifies the set of existing instances of C (and those of each identified supertype of C). A declaration $\mathbf{a} : \mathbf{C}$ in the OWNS list of another class is interpreted as $\mathbf{a} : @C$.

2.2 Specification Examples

Examples of properties which may be expressed in an abstract declarative manner using Z^{++}/RTL formulae are: "m initiates every t seconds, and in the order of its requests":

 $\forall \mathbf{i} : \mathbb{N}_1 \bullet \uparrow (\mathbf{m}(\mathbf{x}), \mathbf{i}+1) = \uparrow (\mathbf{m}(\mathbf{x}), \mathbf{i}) + \mathbf{t}$

 $\uparrow(\mathbf{m}(\mathbf{x}), \mathbf{i})$ denotes the time that the i-th request for invocation of $\mathbf{m}(\mathbf{x})$ received by the current object begins to execute, whilst $\dagger(\uparrow \mathbf{m}(\mathbf{x}), \mathbf{i})$ denotes the i-th time of the form $\uparrow(\mathbf{m}(\mathbf{x}), \mathbf{j})$. These may be different if invocations are not initiated in the order that they are received by an object. The 'shortest job first' protocol is stated:

$$\begin{array}{l} \forall \, \mathbf{i}, \mathbf{j} : \mathbb{N}_1 \mid \rightarrow (\mathbf{m}, \mathbf{i}) \leq \uparrow(\mathbf{m}, \mathbf{j}) \bullet \\ f((\mathbf{m}, \mathbf{i}).\mathbf{x}) < f((\mathbf{m}, \mathbf{j}).\mathbf{x}) \ \Rightarrow \ \uparrow(\mathbf{m}, \mathbf{i}) \leq \uparrow(\mathbf{m}, \mathbf{j}) \end{array}$$

That is, if (m, j) has not already started execution by the time the request for (m, i) arrives, then (m, i) will be started first if it has a smaller value of some priority assigning function f on parameter(s) x.

Liveness and fairness constraints can also be stated. An example of a complete (but highly abstract) Z^{++} class is a binary semaphore:

```
CLASS Semaphore

OPERATIONS

signal: →;

release: →

HISTORY

#fin(signal) ≥ #act(release) ∧

#fin(release) + 1 ≥ #act(signal)

END CLASS
```

self_mutex({signal, release}) and mutex({signal, release}) follow from the remainder of the history constraint, which is implicitly quantified by a \Box^{τ} operator.

3 Z⁺⁺/RTL Logic

3.1 Logic

Events

- For each method m of C: $\uparrow m(e)$, $\downarrow m(e)$, $\rightarrow m(e)$ for $e \in IN$, denoting the initiation, termination and the arrival of a request at the object class, respectively, of an invocation instance of m(e);
- $-\theta :=$ true, $\theta :=$ false for a predicate θ without modal operators, which denote the events of this predicate becoming true or false, respectively.

These events, together with events of the form \leftarrow (n(x1), a) for a : S a supplier object to C, and n a method of S (the sending of a request for a to execute n(x1)), are collectively referred to as **BasicEvent**_C. The complete set of events of C also include the following:

- $-\uparrow(\mathbf{n}(\mathbf{x}1),\mathbf{a}),\downarrow(\mathbf{n}(\mathbf{x}1),\mathbf{a}),\rightarrow(\mathbf{n}(\mathbf{x}1),\mathbf{a})$ where \mathbf{a} and \mathbf{n} are as above;
- $\leftarrow (\mathbf{n}(\mathbf{x}1), \mathbf{a}, \mathbf{b})$ where both \mathbf{a} and \mathbf{b} are supplier objects to \mathbf{C} , $\mathbf{a} : \mathbf{S}$, and \mathbf{n} is a method of \mathbf{S} .

The complete collection of events of C is denoted by Event_C.

Terms For a given class **C**, the following terms can occur in the formulae of its RTL language:

- 1. variables $\mathbf{v_i}$: $\mathbf{i} \in \mathbb{N}$ only variables can be quantified over;
- 2. attributes of the class, its ancestors and supertypes;
- 3. $f(e_1, ..., e_n)$ for an n-ary function symbol f and terms $e_1, ..., e_n$, and other Z expressions in terms and schema texts;
- 4. $\dagger e$ where e is an event occurrence (E, i), where E is in Event_C the time of the i-th occurrence of E, $i : \mathbb{N}_1$;
- 5. Op(m(e), i) where $m \in \underline{methods}(C)$, e in the input type of $m, i : \mathbb{N}_1$ and $Op \in \{\uparrow, \downarrow, \rightarrow\}$, and $\leftarrow ((m(e), a), i)$ for a supplier object a : D and method m of D;
- 6. $e \otimes t$ and $\bigcirc e$ where e is a term, t a time-valued term the value of e at t and at the next method initiation, respectively;
- 7. #act(m(e)), #fin(m(e)), #req(m(e)), #req(m), #fin(m), #act(m) for $m \in \underline{methods}(C)$;
- 8. self.

#req(m) is the number of requests for m received by the current object up to the present time, #act(m) the number of initiations of execution of m and #fin(m) the number of terminations of m.

Time-valued terms are arithmetic combinations of terms of the form 4 or 5 and elements of N. The time domain satisfies the axioms of the set of non-negative elements of a totally ordered topological ring, with addition operation + and unit 0, and multiplication operation * with unit 1. Thus $\mathbb{N} \subseteq \text{TIME}$ can be assumed. Quantification over TIME is not allowed.

Relativised versions #act(m(e), a), etc of event counters for suppliers a : D to C are also included, as are attributes a.att of such suppliers.

In addition, method names m from supertypes D of C can be used in its events. They will be interpreted (if unambiguous) as $\phi(\mathbf{m})$ in the semantics of the language, where $\mathbf{D} \sqsubseteq \phi_{\mathbf{R}} \mathbf{C}$ is asserted in the specification.

 $\overline{\mathbf{A}}$ and $*_{\mathbf{A}}$ can be referred to, for any class \mathbf{A} in the specification. This enables control over object sharing and aliasing.

Formulae For any class C the following are the formulae in its RTL language.

- 1. $P(e_1, \ldots, e_n)$ for an n-ary predicate symbol P and terms e_1, \ldots, e_n ;
- 2. $\phi \land \psi, \phi \lor \psi, \phi \Rightarrow \psi, \neg \phi$ for formulae ϕ and ψ ;
- 3. $\phi \odot \mathbf{t}$ for formulae ϕ and time-valued terms $\mathbf{t} \phi$ holds at time \mathbf{t} ;
- 4. \forall **SD** ϕ , \exists **SD** ϕ for declarations **SD** and formulae ϕ ;
- 5. $\Box^{\tau}\theta$, $\Box\theta$ and $\bigcirc\theta$ for formulae θ ;
- 6. $\diamond^{\tau} \theta$, $\diamond \theta$ for formulae θ ;
- 7. enabled(m) and enabled(m(e)) for methods m, e in the input type of m, and enabled(m, a) for suppliers a of C.

 $\Box \psi$ denotes that ψ holds at all times of the form $\uparrow(\mathbf{m}, \mathbf{i})$ for \mathbf{m} a method of **C** which are at or greater than the present time. $\bigcirc \psi$ denotes that ψ holds at the next time of the form $\uparrow(\mathbf{m}, \mathbf{i})$ (if there is any). In contrast \Box^{τ} and \diamond^{τ} refer to all present and future times.

Axioms Only selected axioms will be presented here. A full list, together with details of the semantics and a proof of soundness, is contained in [7]. The axioms include those of classical predicate logic in this language, and the axioms of the Z mathematical toolkit. Of particular importance is a *frame* axiom which asserts that attributes of C can only change in value if a method of C is currently executing – it is a form of *locality* property in the sense of [4].

The ACTIONS specification of a method m is interpreted as:

$$\begin{aligned} \mathbf{(xi)}: \ \forall \mathbf{e}: \mathbf{IN}; \ \mathbf{i}: \mathbb{N}_1 \bullet \ \mathbf{Pre}_{\mathbf{m},\mathbf{C}} \odot \uparrow (\mathbf{m}(\mathbf{e}),\mathbf{i}) \Rightarrow \\ \mathbf{Def}_{\mathbf{m},\mathbf{C}}[\mathbf{v} \circledast \downarrow (\mathbf{m}(\mathbf{e}),\mathbf{i})/\mathbf{v}'] \odot \uparrow (\mathbf{m}(\mathbf{e}),\mathbf{i}) \end{aligned}$$

where the appropriate versions $(\mathbf{m}(\mathbf{e}), \mathbf{i}) \cdot \mathbf{x}_{\mathbf{j}}$ of formal input or output parameters are used in $\mathbf{Pre}_{\mathbf{m},\mathbf{C}}$ or $\mathbf{Def}_{\mathbf{m},\mathbf{C}}$.

Axioms of linear temporal logic (LTL) [15] hold in this formalism, including the induction scheme:

$$\phi \odot \min(\{\dagger(\uparrow \mathbf{m}_1, 1), \dots, \dagger(\uparrow \mathbf{m}_n, 1)\}) \land \Box(\phi \Rightarrow \bigcirc \phi) \Rightarrow \Box \phi$$

which holds for each LTL formula ϕ , where <u>methods</u>(C) = {m₁,...,m_n}.

In [7] it is shown that a version of Manna-Pnueli logic [14] is provable from Z^{++}/RTL , and that Z^{++}/RTL is conservative over Manna-Pnueli logic.

Abbreviations $\underline{methods}(\mathbf{C})$ abbreviates the set of methods of a class \mathbf{C} , including inherited methods.

#active(m) abbreviates #act(m) - #fin(m), the number of currently active instances of m. #waiting(m) abbreviates #req(m) - #act(m), the number of instances of m awaiting execution.

delay(m,i) abbreviates $\uparrow(m,i) - \rightarrow(m,i)$. duration(m,i) abbreviates $\downarrow(m,i) - \uparrow(m,i)$. mutex($\{m_1, \ldots, m_n\}$) abbreviates the assertion

$$#active(\mathbf{m}_1) = \sum_{i=1}^{n} #active(\mathbf{m}_i) \lor \ldots \lor #active(\mathbf{m}_n) = \sum_{i=1}^{n} #active(\mathbf{m}_i)$$

self_mutex($\{m_1, \ldots, m_n\}$) abbreviates $\#active(m_1) \leq 1 \land \ldots \land \#active(m_n) \leq 1$.

 $\underline{\mathbf{m}}$ abbreviates #active $(\mathbf{m}) > 0$.

fires(t,i) denotes \neg (false \odot \uparrow (t,i)), that is, the object exists at this time point.

A durative method **m** is a method which satisfies durative(**m**): $\forall \mathbf{i} : \mathbb{N}_1 \bullet \downarrow(\mathbf{m}, \mathbf{i}) > \uparrow(\mathbf{m}, \mathbf{i})$. A durative class is a class all of whose methods are durative. For such a class, the property $\forall \mathbf{i} : \mathbb{N}_1 \bullet \mathbf{m} \odot \uparrow(\mathbf{m}, \mathbf{i})$ holds for each method **m**.

3.2 Semantics

A model Ω of a Z⁺⁺/RTL specification **S** consists of a family (@C)_{C∈classes_of}(**S**) of countably infinite sets of object references for each class **C** of **S**, and a family $(\Omega_{\mathbf{C}})_{\mathbf{C}\in \underline{classes_of}(\mathbf{S})}$ with typing $\Omega_{\mathbf{C}} : @C \to \mathbf{Object}_{\mathbf{C}}$ where $\mathbf{Object}_{\mathbf{C}}$ is the set of pairs $\beta = (\alpha, \sigma)$ with the type

$$(\alpha, \sigma) : (\mathbf{BasicEvent}_{\mathbf{C}} * \mathbb{N}_1 \rightarrow \mathbf{TIME}) \times (\mathbf{TIME} \twoheadrightarrow \mathbf{State}_{\mathbf{C}}^{\beta})$$

which satisfy a set of conditions corresponding to basic properties of events. State^{β}_C is State_C with each declaration att : T of C replaced by att : T^{β}.

 α assigns a time (not necessarily in the lifetime of the object) to each occurrence of an event of C, and σ gives the state of the object at each time point in its history. From these the value of terms and formulae of $\mathcal{L}_{\mathbf{C}}$ can be computed at each time point.

4 Development Process

The development process envisaged for the use of the language for real time systems is as follows:

- 1. requirements capture and analysis of the problem, using a structured method suitable for real time problems, such as OMT [16];
- 2. formalisation of structured method notations in Z^{++}/RTL , using systematic processes for the translation of object classes in OMT into Z^{++} classes (similar to the processes for B AMN described in [11]) and statecharts into history constraints. The abstract declarative nature of RTL allows fairness, liveness and safety constraints to be stated in an implementation-independent manner;
- refinement of Z⁺⁺/RTL specifications into implementation-oriented classes, making use of reusable specified components and the code of these components;
- 4. *implementation* of classes using classes which contain procedural code, with timing information derived from a particular execution environment being used to prove the final refinement step.

In the following sections stages 2, 3 and 4 will be illustrated using a small example. Larger applications are given in [7, 8].

4.1 Formalisation of OMT Analysis Models

Integrating formal and structured methods has a number of advantages as a development approach: it can make use of the complementary strengths of these two techniques, and it can make use of existing software engineering expertise, rather than attempting to replace it. A recent survey on the use of formal methods in industry reported that 31% of those companies using formal methods were using them in conjunction with structured methods [2].

Object Models There are two main models which form the input to the formalisation process. The first is the *Object Model* which describes the entities involved in the system, and their attributes, operations and the relationships between them (including inheritance or subtyping). An object model describing a railway station consisting of a set of track sections is shown in Fig. 1.

These models are used to build an initial outline specification, which will later be enhanced by consideration of the dynamic model, and by the addition of semantic detail which could not be expressed in the structured models.

The formalisation process is as follows:

- 1. For each entity C in the object model, create a Z^{++} object class C;
- 2. Each attribute of C becomes an attribute of C, with corresponding type;



Fig.1. Object Model of Station

- 3. Each association r between entities C and D is examined to determine if both directions of the association are required in the final system. The required directions are then formalised as attributes \mathbf{r}_1 of C, of type D (in the case of a many-one or one-one association from C to D), or of type $\mathbb{F}(\mathbf{D})$ or seq(D) (in the case of an unordered or ordered many-many or one-many association from C to D, respectively). Similarly for the inverse map \mathbf{r}_2 from D to C, if this is required;
- If D inherits from C, then the clause EXTENDS C is placed in the header of D. Conformant subtypes are expressed via suitable ⊆ assertions;
- 5. Operations are translated into outline specifications of operations, with however all signatures completed;
- 6. Formalisation of all assertions on an object model can be attempted, using the INVARIANT and HISTORY components of a class.

Dynamic Model For reactive systems the dynamic model, based upon Harel statecharts, is the most significant analysis model. Here we will consider extensions of statecharts to include time bounds on transitions. The formalisation process for dynamic models is:

- 1. Condition triggered transitions and anonymous (automatic) transitions are formalised as operations which are internally invoked, whilst event-triggered transitions are formalised as operations which are invokable from other objects;
- 2. states are formalised as elements of an enumerated set, and an attribute of this type is defined to record the current state. Methods formalising a transition modify this variable appropriately;
- 3. time bounds [l, u] on transition t are expressed via the formula $\forall i : \mathbb{N}_1 \bullet \text{fires}(t, i) \Rightarrow l \leq \text{delay}(t, i) \leq u;$

4. if t has source state S1 and destination state S2 (assumed distinct) and guard condition cond, then:

```
\begin{array}{ll} (\textbf{enabled}(\textbf{t}) \equiv (\textbf{state} = \textbf{S1}) \land \textbf{cond}) & \land \\ \forall \textbf{i} : \mathbb{N}_1 \bullet \exists \textbf{j}, \textbf{k1}, \textbf{k2} : \mathbb{N}_1 \bullet \\ & \dagger ((\textbf{state} = \textbf{S1}) \land \textbf{cond} := \textbf{true}, \textbf{j}) = \rightarrow (\textbf{t}, \textbf{i}) \land \\ & ((\textbf{state} = \textbf{S1}) \land \textbf{cond}) \odot \uparrow (\textbf{t}, \textbf{i}) \land \\ & \downarrow (\textbf{t}, \textbf{i}) = \dagger ((\textbf{state} = \textbf{S1}) := \textbf{false}, \textbf{k1}) \land \\ & \downarrow (\textbf{t}, \textbf{i}) = \dagger ((\textbf{state} = \textbf{S2}) := \textbf{true}, \textbf{k2}) \end{array}
```

If t is durative (in particular if $S1 \neq S2$) then the LTL properties

 $\Box(\underline{\mathbf{t}} \Rightarrow \bigcirc (\mathbf{state} = \mathbf{S}2)) \qquad \Box(\underline{\mathbf{t}} \Rightarrow \mathbf{state} = \mathbf{S}1)$

can be derived (if t has no other source or destination);

5. if t has an associated action act on supplier object a, then:

 $\forall \mathbf{i} : \mathbb{N}_1 \bullet \exists \mathbf{k}3 : \mathbb{N}_1 \bullet \uparrow (\mathbf{t}, \mathbf{i}) = \leftarrow ((\mathbf{act}, \mathbf{a}), \mathbf{k}3)$

6. event triggered transitions are formalised in the same way, however clause 4 is replaced by:

$$\begin{array}{l} \mathbf{enabled}(\mathbf{t}(\mathbf{p})) \equiv (\mathbf{state} = \mathbf{S1}) \land \\ \forall \mathbf{i} : \mathbb{N}_1 \bullet \exists \mathbf{j} : \mathbb{N}_1 \bullet \\ (\mathbf{state} = \mathbf{S1}) \odot \dagger (\mathbf{event}(\mathbf{p}), \mathbf{j}) \land \\ \dagger (\mathbf{event}(\mathbf{p}), \mathbf{j}) = \rightarrow (\mathbf{t}(\mathbf{p}), \mathbf{i}) \land \\ (\mathbf{state} = \mathbf{S1}) \odot \dagger (\mathbf{t}(\mathbf{p}), \mathbf{i}) \end{array}$$

Each class corresponding to a statechart is mutex and self-mutex. In addition there are liveness constraints asserting that any non-terminal state must eventually be exited, and constraints asserting that states can only become true as a result of a transition into them. If there are no self-transitions on S1 then the assertion \Box (state = S1 \Rightarrow $\underline{t_1} \lor \ldots \lor \underline{t_n}$) states that S1 can only be exited via transitions t_1, \ldots, t_n .

For example, consider the statechart of track sections shown in Fig. 2.

A track section is a defined contiguous segment of track which can be occupied by at most one train at any time. In addition, it may be closed (eg, for engineering work) so that no trains may enter the section. A train will take a minimum of 60 seconds to clear a track section, and a closed track section will be closed for a minimum of 120 seconds. The corresponding outline class, with a state variable tstate : TState is:

```
CLASS TrackSection
TYPES
TState ::= closed | free | blocked
OWNS
tstate : TState
```



Fig. 2. Statechart of TrackSection

```
OPERATIONS
  init: \rightarrow;
  train_arrives : \rightarrow;
  train_departs : \rightarrow;
  open: \rightarrow;
  close: \rightarrow
ACTIONS
  init ==>
             tstate' = free;
  train_arrives ==>
             tstate' = blocked;
  train_departs
                       ==>
             tstate' = free;
            ==>
  open
             tstate' = free;
  close
            ==>
             tstate' = closed
HISTORY
  mutex(\{ init, train_arrives, train_departs, open, close \}) \land
  self_mutex({ init, train_arrives, train_departs, open, close }) \land
  \Box(\underline{\text{train\_arrives}} \Rightarrow \text{tstate} = \text{free}) \land
  \Box(train_departs \Rightarrow tstate = blocked) \land
  \Box(open \Rightarrow tstate = closed) \land
  \Box(\underline{\text{close}} \Rightarrow \text{tstate} = \text{free}) \land
  \forall i: \mathbb{N}_1 • fires(train_departs, i) \Rightarrow 60 \leq delay(train_departs, i) \land
```

```
\forall i: \mathbb{N}_1 \bullet fires(open, i) \Rightarrow 120 \leq delay(open, i) END CLASS
```

There are additional logical properties which can be derived from the statechart. More detail on the formalisation process, including the treatment of entry and exit actions, activities within a state, nested statecharts and AND composition of statecharts is given in [8]. The precision of the translation is such that statecharts could be explicitly presented in the HISTORY clause of a class in place of (some) temporal assertions.

4.2 Refinement

Refinement in Z^{++}/RTL corresponds to theory extension [6, 7]. That is, all specified dynamic properties of a class C should be provable in any refinement D of C. The notation for refinement of C by D, via a renaming ϕ of the methods of C, and a data refinement relation R on the combined states is $C \sqsubseteq \phi_{\mathbf{R}} \mathbf{D}$.

Subtyping of classes is equated with class refinement, in contrast to [12], although the roles of subtyping and refinement within the development process are clearly distinguished. Alternative concepts of refinement, related to *operational compatibility* [3], and strictly stronger than subtyping, are discussed in [7].

4.3 Implementation

In the final refinement of a subsystem, a restricted language is used, in which methods are defined using constructs corresponding to procedural code structures, and in which types are restricted to be arrays, strings, object reference types or scalars.

This language preserves refinement in the sense that if a class **D** is a client of a class **C**, and **C**₁ refines **C**, then substituting **C**₁ for **C** in **D** to produce a class **D**₁ implies that **D**₁ refines **D**. This supports compositional (separate) development of subsystems. It is also direct to implement such classes in an object-oriented programming language such as C++. Timing specifications for such implementations can be derived from the timing specifications of the hardware which is the ultimate destination of the compilation process.

The RTL language is extended by times $\uparrow(s, i)$, $\downarrow(s, i)$ and $\rightarrow(s, i)$ for each statement occurrence s of the implementation language in C. These denote respectively start, end, and request times for the *i*-th invocation of s within an instance of C.

These times are assumed in general to be non-negative real numbers. The approach of [5] is used: For each primitive operation op, which is one of: gaddr (get an address), stor (store a value in an address), eval(e) (evaluate expression e), decl(T) (define storage for a variable of type T), brt, brtloop (branch on true), brf, brfloop (branch on false), there are corresponding sets T(op) of possible durations for this operation.

T(op) should always be a small finite set, thus allowing reasoning by case analysis. The definition of T will vary between destination processor architectures. An operation $\dot{+}$ performs addition of such sets, ie: $S \dot{+} T = \{s + t \mid s \in S \land t \in T\}$. An example of the timing rules is:

Assignment The behaviour of assignment is given by:

That is, there is no delay in executing the assignment after its request, and its duration is a sum of possible durations for its constituent steps.

5 Development Using Semaphores

A general approach for the refinement of classes with non-trivial synchronisation constraints is to utilise classes, such as semaphores, which provide specific synchronisation facilities. That is, a class C with synchronisation requirements expressed in its HISTORY clause will be refined by a class C_1 which has a supplier class S whose properties can be used to prove the requirements of C. C_1 itself may not need to contain any synchronisation mechanisms.

An example is the case where it is required that a particular method \mathbf{m} is self-mutex (Fig. 3). Synchronisation constraints of \mathbf{C} are implemented using a

```
CLASS C_1
                                         OWNS
CLASS C
                                           s: Semaphore;
OWNS
                                           ...
                                         OPERATIONS
OPERATIONS
                                           m: X \rightarrow Y;
  \mathbf{m}: \mathbf{X} \rightarrow
                                         ACTIONS
ACTIONS
                                           mхy
                                                      ==>
                     Def<sub>m</sub>;
  mxy
                                                      BEGIN
                                                         s.signal;
HISTORY
                                                         Code;
  self_mutex({ m })
                                                         s.release
END CLASS
                                                      END;
                                         END CLASS
```

Fig. 3. Refinement Using Semaphores

Semaphore instance, where Code implements the state transitions defined in Def_m . The refinement is formally provable because:

$$\begin{aligned} \#active(\mathbf{m}) &= \#act(\mathbf{m}) - \#fin(\mathbf{m}) \\ &= \#act(signal, s) - \#fin(release, s) \le 1 \end{aligned}$$

from the history constraint of Semaphore.

Other examples are given in [8], including a development of the dining philosophers problem using this approach.

6 Refinement Using Synchronisation Code

An alternative refinement route, which allows synchronisation and fairness requirements expressed in a Z^{++} specification to be discharged, makes use of the concept of synchronisation code [13]. This involves the definition of state transformations on synchronisation variables which are performed at event occurrences. An event is either the creation of an object, a reception of a request for a method execution, the start of an execution of a method, or the termination of execution of a method. Thus the approach is consistent with the Z^{++}/RTL formalism, and with statecharts.

Synchronisation variables and functions over these are then used to control the permission to execute methods via permission guards.

The refinement of a class (at the end of a data and procedural refinement process) into a C++ class with synchronisation code can be formally checked by an induction over events. This induction will be that the abstract declarative history constraint of the Z^{++} class is always true at each event time, with respect to the translation between Z^{++} and C++ variables.

As an example, consider the specification that requests of \mathbf{m} are served in a first-come, first-served manner (FCFS):

An appropriate extended C++ class is (following [13]):

```
class C{
  m() {}
synchronisation
  int clk;
  int arr_time local to m;
  start(C) --> clk = 0;
  arrival(m) --> this_inv.arr_time = clk++;
  m: there_is_no(p in waiting(m): p.arr_time < this_inv.arr_time);
}</pre>
```

The final clause here asserts that the **this_inv** invocation instance of **m** cannot initiate execution unless the priority condition holds for this instance.

We can relate these two versions of C as follows. The concrete expression this_inv.arr_time refers to a local (synchronisation) variable of a specific instance (m, i) of an invocation of m. This can be formalised as (m, i).arr_time. From the code it follows that: (m, i).arr_time = $clk \circledast \rightarrow (m, i)$.

By induction on events it can be shown that this value is also

$$\textbf{previous_requests}(\textbf{m}, \rightarrow(\textbf{m}, \textbf{i})) \; = \; \#\{\textbf{j} : \mathbb{N}_1 \mid \rightarrow(\textbf{m}, \textbf{j}) < \rightarrow(\textbf{m}, \textbf{i})\}$$

(The only relevant events are those listed in the concrete class, and it is assumed that object creation occurs before any (m, i) invocation.)

The set waiting(m) of outstanding invocation instances of m has the formal counterpart waiting_instances(m, t) = $\{i : \mathbb{N}_1 \mid \rightarrow(m, i) \leq t \land \neg (\uparrow(m, i) \leq t)\}$ for each time-valued term t.

Therefore the concrete guard has corresponding formal permission requirement

```
 \begin{array}{l} \forall \mathbf{i} : \mathbb{N}_1 \bullet \\ (\neg \ \exists \ \mathbf{pind} : \mathbb{N}_1 \ | \ \mathbf{pind} \in \mathbf{waiting\_instances}(\mathbf{m}, \uparrow(\mathbf{m}, \mathbf{i})) \bullet \\ & \mathbf{previous\_requests}(\mathbf{m}, \rightarrow(\mathbf{m}, \mathbf{pind})) < \\ & \mathbf{previous\_requests}(\mathbf{m}, \rightarrow(\mathbf{m}, \mathbf{i}))) \circledast \uparrow(\mathbf{m}, \mathbf{i}) \end{array}
```

Given this, it is impossible for there to be pind $\langle i with \rightarrow (m, pind) \rangle \langle \rightarrow (m, i) but \uparrow (m, i) \langle \uparrow (m, pind) \rangle$.

Assume otherwise. Then pind \in waiting instances $(m, \uparrow(m, i))$. But

 $previous_requests(m, \rightarrow (m, pind)) < previous_requests(m, \rightarrow (m, i))$

since **pind** is a member of the second set and not the first (and **previous_requests**(m, t) is monotonically increasing with t), contradicting the guard for m.

Tools

Animation and proof tools have been developed for this formalism. Animation is based upon the checking of proposed scenarios expressed as sequences of events. The tool allows a sequence of events and corresponding times to be incrementally constructed, checking permission guards, duration and delay constraints, synchronisation constraints involving event counters, and mutual exclusion properties. Post-states are generated from pre-states, with user interaction being required if there is non-determinism in the method specification [9].

Conclusions

This paper has detailed development processes for real time system development which integrate structured and formal methods in an object-oriented framework, and which combine the benefits of these methods. The proposed formalism provides many of the facilities of real time formalisms such as interval logic [1] or RTTL [14] without excessive notational overhead.

References

- J. F. Allen: Maintaining knowledge about temporal intervals, CACM, 26(11):832– 843, November 1983.
- S. Austin, G. I. Parkin: Formal Methods: A Survey, National Physical Laboratory, Queens Road, Teddington, Middlesex, TW11 0LW, March 1993.
- R. Duke, P. King P. G. Smith: Formalising Behavioural Compatibility for Reactive Object-oriented Systems, in Proc 14th Australian Compt. Sci. Conf. (ACSC-14), 1991.
- J. Fiadeiro, T. Maibaum: Sometimes "Tomorrow" is "Sometime", in Temporal Logic, D. M. Gabbay and H. J. Ohlbach (editors), LNAI 827, Springer-Verlag 1994, 48-66.
- 5. C. Fidge: Proof Obligations for Real-Time Refinement, Proceedings of 6th Refinement Workshop, Springer-Verlag Workshops in Computing, 1994.
- 6. K. Lano: Refinement in Object-oriented Specification Languages, Proceedings of 6th Refinement Workshop, Springer-Verlag Workshops in Computing, 1994.
- K. Lano: Formal Object-oriented Specification of Real Time Systems, Dept. of Computing, Imperial College, 1994.
- 8. K. Lano: Software Specification and Development in Z⁺⁺, to appear in The Z Handbook, J. Bowen and M. Hinchey (eds.), McGraw-Hill 1995.
- 9. K. Lano: Reasoning Techniques in VDM⁺⁺, AFRODITE project report AFRO/IC/KL/RT/V1, Dept. of Computing, Imperial College, 1994.
- K. Lano, H. Haughton: Object-oriented Specification Case Studies, Prentice Hall, 1993.
- 11. K. Lano, H. Haughton: Improving the Process of System Specification and Refinement in B, Proceedings of 6th Refinement Workshop, Springer-Verlag Workshops in Computing, 1994.
- 12. B. Liskov, J. Wing: Family Values: A Behavioral Notion of Subtyping, School of Computer Science, Carnegie Mellon University, report CMU-CS-93-187, 1993.
- 13. C. McHale, S. Baker, B. Walsh, A. Donnelly: Synchronisation Variables, Amadeus Project report TCD-CS-94-01, University of Dublin, 1994.
- 14. J. S. Ostroff: Temporal Logic for Real-Time Systems, John Wiley, 1989.
- 15. A. Pnueli: Applications of Temporal Logic to the Specification and Verification of Reactive Systems: A Survey of Current Trends, Current Trends in Concurrency, de Bakker J., de Roever W.-P., Rozenberg G. (Eds), Springer-Verlag Lecture Notes in Computer Science, Vol. 224, 1986.
- 16. J. Rumbaugh, M. Blaha, W. Premerlani, F. Eddy, W. Lorensen, Object-Oriented Modelling and Design, Prentice-Hall International, 1991.