

A Timeband Framework for Modelling Real-Time Systems

Alan Burns · Ian J. Hayes

Abstract Complex real-time systems, must integrate physical processes with digital control, human operation and organizational structures. New scientific foundations are required for specifying, designing and implementing these systems. One key challenge is to cope with the wide range of time scales and dynamics inherent in such systems. To exploit the unique properties of time, with the aim of producing more dependable computer-based systems it is desirable to explicitly identify distinct time bands in which the system is situated. Such a framework enables the temporal properties and associated dynamic behaviour of existing systems to be described and the requirements for new or modified systems to be specified. A system model based on a finite set of distinct time bands is motivated and developed in the paper.

1 Introduction

The construction of large socio-technical real-time systems, such as those envisaged in cyber-physical applications, imposes a number of significant challenges, both technical and organisational. Their complexity makes all stages of their development (requirements analysis, specification, design, implementation, deployment and maintenance/evolution) subject to failure and costly re-working. Even the production of an unambiguous behavioural description of an existing system is far from straightforward.

One characteristic of these computer-based systems is that they are required to function at many different time scales (from microseconds or less to days or more). Time is clearly a crucial notion in the specification (or behavioural description) of computer-based systems, but it is usually represented, in modelling schemes for example, as a single flat physical phenomenon. Such an abstraction fails to support the structural properties of the system, forces different temporal notions on to the same

A. Burns
Department of Computer Science, University of York, UK.

I.J. Hayes
The University of Queensland, School of Information Technology and Electrical Engineering, Australia.

flat description, and fails to support the separation of concerns that the different time scales of the system facilitate. Even with a single time scale, system architects seem to have great difficulty in specifying temporal properties in anything other than very concrete implementation-level terms. But just as the functional properties of a system can be modelled at different levels of abstraction or detail, so too should its temporal properties be representable in different, but provably consistent, time scales.

To make better use of ‘time’, with the aim of producing more dependable computer-based systems, we propose a framework that explicitly identifies a number of distinct *time bands* in which the system under study is situated. The framework enables the temporal properties of existing systems to be described and the requirements for new or modified systems to be specified.

In the following section we first motivate the main notions and properties of the timeband framework. Then, in Section 3, an abstract model of timebands is presented, and in Section 4 is extended to describe state. Section 5 gives a brief summary of the model. The model is, in itself, not intended to be a complete semantic description. That is achieved by ‘embedding’ the model in a parent notation/logic. The focus of this paper is, however, the framework itself. Section 6 introduces some notation to allow specification of properties in the timeband framework. A short example of the use of the framework is presented in Section 7 and conclusions are covered in Section 8.

2 Motivation

A large real-time system exhibits dynamic behaviour on many different levels. The computational components have circuits that have nanosecond speeds, faster electronic subcomponents and slower functional units. Communication on a fast bus is at the microsecond level but may be tens of milliseconds on slow or wide-area media. Human time scales move from the 1ms neuron firing time to simple cognitive actions that range from 100ms to 10 seconds or more. Higher rational actions take minutes and even hours. Indeed it takes on the order of 1000 hours to become an expert at a skilled task, such as flying a plane [22] and the development of highly skilful behaviour may take many years. At the organisational and social level, time scales range from a few minutes, through days, months and even years. Perhaps for some environmentally sensitive systems, consequences of failure may endure for centuries. To move from nanoseconds to centuries requires a framework with considerable descriptive and analytical power.

The concept of timebands comes from an detailed study of existing computer-based systems¹ and their requirements (eg.[13]), ethnographical studies (eg. [5,2]), the work of Newell [18] in his attempts to describe human cognition, work on system structure such as that of Simon [23], reports on system failures (eg. Columbus [4]), studies from areas such as the psychology and sociology of time [9,10,21,17,3], formalisms such as the teleo-reactive programming model [19,20] and Statecharts that require instantaneous state changes, and the few examples of modelling work

¹ As part of the Dependability Interdisciplinary Research Collaboration – DIRC, see web.

that do attempt to consider more than one time scale within a system (eg. Corsetti *et al* [8, 7]).

As the concept of ‘now’ (present moment) seems to be fundamental to our reasoning about time, it follows that notions such as ‘instantaneous’, ‘simultaneous’ and ‘immediate’ are natural ones to use in specifying temporal properties. As Bergadaà [3] states in his work on temporal framework:

The present time could be made of moments that enable the allocation of time to different activities. They could also be made of duration in which the activity will take the time necessary for its completion.

From such observations and the literature noted above we distil the following properties that we identify as being of relevance to the modelling of complex real-time systems.

- The dynamics of a system (how quickly things change) are central to understanding its behaviour.
- Systems clearly operate at many different granularities (of time), ie. there are different abstract views of the dynamics of the system.
- It is useful to consider certain actions (events) as atomic and instantaneous (whilst allowing them to have internal state and behaviour at a more detailed level of description).
- It is useful to consider two or more events as occurring simultaneously (instantaneously), or the response to some event being immediate (whilst allowing them to be separated in time at a more detailed level of description).
- The order (but not necessarily the time) at which events occur is important; precedence can give rise to causality.
- The durations of certain actions are important, but the measuring of time must not be overly precise and must allow for tolerance (non-determinacy) in the temporal domain.
- Abstract clocks are useful for relating and coordinating activities, but real clocks are never perfectly reliable or accurate.
- At each level of temporal behaviour it is useful to have access to both continuous and discrete notions of time – controlling actions are typically described using discrete time, the controlled object due to its continuously changing nature often requires dense time for its behavioural description.
- Hierarchical control (cascade control) and hierarchical scheduling (planning) are often observed through the time levels of a system.
- At each level of temporal behaviour similar phenomena is observed (e.g., cyclic/repetitive actions, deadlines-driven actions, synchronous and asynchronous event handling, agreement, coordination, etc.)

Engineers, even of real-time systems, seem to have great difficulty over the use of precise values of time. Why choose an iteration rate of 20ms? – why not 19ms or 21ms? What does a deadline of 15ms actually mean? – would a delay of 10 μ s be significant? This difficulty with temporal quantities is not mirrored in the physical domain where tolerances on lengths, weight etc. are commonly expressed.

In the timebands framework, apparently more natural (and essentially atemporal) notions are available such as ‘immediate’, ‘instantaneous’, ‘simultaneous’ ‘definitely’ and ‘possible’. And durations are first expressed in general terms - for example “this is a minute-level activity” (ie. it will last a few minutes, rather than hours or seconds). Orders of magnitude between rates of change give an initial decomposition of the system. Indeed the framework uses time itself to separate concerns in any architectural description or system specification.

The central notion in the framework is that of a time band that is defined by a granularity (eg. 1 minute) and a precision (eg. 5 seconds). Granularity defines the unit of time of the band; precision bounds the actual duration of an event that is deemed to be instantaneous in this band. To complete this short motivation section the important topics of sampling and rates of change are addressed.

Sampling. Our focus is on embedded real-time systems and hence we cannot avoid issues like sampling of inputs, and discretization of continuous quantities. For example, if it is an error for two proximity sensors to give simultaneous positive inputs, then by placing this requirement for error detection in the minute band (with precision of 5 seconds) the following constraints are derived – assume the two proximity conditions are represented by boolean variables *top* and *bottom*, representing that a controlled gate is at the top or bottom, respectively, of its travel.

- If in any interval of duration five seconds, or more, *top* and *bottom* are permanently true then the error condition *must* be identified.
- If in any interval of duration five seconds, or less, *top* and *bottom* are true for part of the interval then the error condition *can* be identified. Note that the intervals over which *top* and *bottom*, respectively, hold don’t have to be the same, or even overlap.

This dual use of *must* and *can* cannot be eliminated. One may move the requirement between time bands to decrease the value of the precision parameter, but even in the lowest band in the system there is an inevitable non-determinacy because true perfectly simultaneous polling of the two sensors is not possible.

Rate of change. Even though an environmental entity is subject to continuous change it does not mean that all such behaviour must be captured at the lowest possible band – and that this band must have a dense notion of time whilst the others can be discrete. Rather within any band many (perhaps most) entities will be discrete, but some may be continuous. So if the purpose of an automatic ‘plant watering’ system in a greenhouse is to aid the growth of plants in some controlled environment, the rate of growth of the crop per week or day may be significant (but not per second or millisecond).

Consider, for example, the maximum rate of flow of water from a pump. At an upper time band this may be stated as r litres per time unit. But at a lower time band, for a piston-style pump there are two phases: one filling the cylinder, in which there is almost no flow of water out of the pump; and the other emptying the cylinder, during which the rate of flow of water out of the pump is about twice r .

The rate change of a state variable may be uniform between some pairs of time bands, but not between others. For example, with the piston pump the rate of flow is not uniform at the time band that distinguishes the filling and emptying phases of the cylinder, but between a pair of higher bands the rate of flow may be uniform. At a still higher pair of bands, we may be switching the pump on and off to control the rate of flow over a broader time base. Again the rate of flow won't be uniform, but between still higher bands, which don't distinguish the on and off phases, it may emerge to be uniform again.

The motivation for proposing this timeband framework is to simplify the specification of complex systems, improve the dependability of deployed systems and reduce the cost of designing (and redesigning) such systems. It allows dynamic properties to be partitioned but not isolated from each other.

3 Definition of the Timeband Model

From the above considerations, a timeband model has been developed² that is described in this section (with some illustrative small examples). The aim of a timeband model is to be an essential part of any complete system description. It enables the temporal properties of existing systems to be described and the requirements for new or modified systems to be specified. The informal description of the framework is supported by a formal model expressed in the Z notation [24].

The framework is developed in a number of stages that build up the full model. Some examples of how this model can be extended into a language for actual use in specifying systems is then given. The list of topics discussed in this section are: time bands, granularity and precision, events and classes of events, precedence, simultaneous and immediate, activities, mappings between bands, durations, clocks, (nondeterministic) state, change-of-state events, mapping states, and rates of change.

3.1 Time bands

A system is assumed to consist not of a single time dimension but a finite set of *bands*. Each band is represented by a *granularity* (expressed as a unit of time that has meaning within the band, e.g. the minute band) and a *precision* that is a measure of the accuracy of the time frame defined by the band – see definitions below. System activities are placed in some band B if they engage in significant events at the time scale represented by B. They have dynamics that give rise to changes that are observable or meaningful in band B's granularity. So, for example, at the 10 millisecond band, neural circuits are firing, significant computational functions are completing, and an amount of data communication will occur. At the 5 minute band, work shifts are changing, meetings are starting, etc. For any system there will be a highest and lowest band that gives a temporal system boundary – although there will always be

² Initial developments of the framework are described in technical reports [6,5].

the potential for larger and smaller bands. Note that at higher bands the physical system boundary may well be extended to include wider (and slower) entities such as legislative constraints or supply chain changes.

For our formal model, we take the set of time bands (\mathcal{B}) as a primitive type. Both “ \sqsubseteq ” and “ \sqsubset ” are relations between bands, with “ \sqsubseteq ” forming a partial ordering³ on time bands. The type of a relation between bands is given as $\mathcal{B} \leftrightarrow \mathcal{B}$. A relation is equivalent to a set of pairs, ie. $(\mathcal{B} \leftrightarrow \mathcal{B}) = \mathbb{P}(\mathcal{B} \times \mathcal{B})$, where $\mathbb{P}X$ stands for power set of X , (i.e., the set of all subsets of X).

$$\left| \begin{array}{l} - \sqsubseteq - : \text{partial_order}[\mathcal{B}] \\ - \sqsubset - : \mathcal{B} \leftrightarrow \mathcal{B} \end{array} \right. \\ \hline \forall b1, b2 : \mathcal{B} \bullet (b1 \sqsubset b2 \Leftrightarrow b1 \sqsubseteq b2 \wedge b1 \neq b2)$$

For example, we may have that

$$\text{MinuteBand} \sqsubset \text{DayBand} \sqsubset \text{MonthBand}$$

From a focus on some band B adjacent bands are identified ($C \sqsubset B \sqsubset A$). Slower (higher or coarser) bands (e.g. A) can be taken to be unchanging (essentially constant) for issues of concern to B . At the other extreme, behaviours in the faster (lower or finer) bands (e.g. C) are assumed to be instantaneous. The actual differences in granularity between A , B and C are not precisely defined (and indeed may depend on the bands themselves) but will typically be in the range 1/10th to 1/100th. When bands map on to hierarchies (structural or control) then activities in band A can be seen to constrain the dynamics of band B , whereas those in C enable B to proceed in a timely fashion. The ability to relate behaviour at different time bands is one of the main properties of the framework.

As an example, consider a university lecture course. Here there are immediately four bands to identify. The year band in which new courses and curriculum are planned, the weekly band in which lectures are scheduled, the minute band that allows each lecture to be structured, and the second band that can capture various interactions with the available technical support (eg. laptop response). Whilst giving a lecture, one can assume that the curriculum is stable (unchanging) and that the laptop reacts instantaneously to slide change requests.

Systems that don’t respect some form of time band structure can become extremely complex and difficult to comprehend, e.g., changing the course syllabus while lecturing is likely to lead to great confusion.

It is important to emphasise that the full behaviour of a system is not obtained by refining down to the lowest band or by projecting emergent behaviours up to the highest band. Rather it is the amalgamation of all band descriptions – all have behaviours that may be needed in any assertion about the system as a whole.

³ See Appendix A.

3.2 Granularity and precision

For each band its *granularity*, representing the unit of time in that band, and *precision*, representing the measure of accuracy of events within that band. They must both be expressed relative to a lower band. For example, the granularity of the *MonthBand* with respect to the *DayBand* may have a granularity defined as follows:

$$\text{Granularity}(\text{MonthBand}, \text{DayBand}) = \{28, 29, 30, 31\}$$

and the granularity of the *DayBand* with respect to the *MinuteBand* is defined as follows:

$$\text{Granularity}(\text{DayBand}, \text{MinuteBand}) = \{1440\},$$

because there are $24 * 60 = 1440$ minutes in a day. Hence the granularity of the *MonthBand* with respect to the *DayBand* is

$$\begin{aligned} \text{Granularity}(\text{MonthBand}, \text{MinuteBand}) = \\ \{28 * 1440, 29 * 1440, 30 * 1440, 31 * 1440\}. \end{aligned}$$

Note that this set is not contiguous. For ease of presentation we assume that standard units such as minutes, milliseconds etc are well defined. However not all time scales will give rise to time bands.

For a band b_1 , its granularity will be defined with respect to all lower bands; hence the domain of the granularity function is all pairs of bands (b_1, b_2) , such that b_2 is lower than b_1 . If b_1 is related to a lower band b_2 , and b_2 to b_3 , then the granularity of b_1 with respect to b_3 is the composition of the granularities of b_1 with respect to b_2 and b_2 with respect to b_3 . The set \mathbb{N} is the natural numbers and \mathbb{N}_1 is the non-zero natural numbers. *Granularity* is a partial function (\leftrightarrow) from pairs of time bands to a non-empty set (\mathbb{P}_1) of non-zero natural numbers.

$$\left| \begin{array}{l} \text{Granularity} : (\mathcal{B} \times \mathcal{B}) \leftrightarrow \mathbb{P}_1 \mathbb{N}_1 \\ \hline \text{dom}(\text{Granularity}) = \{b_1, b_2 : \mathcal{B} \mid b_2 \sqsubset b_1\} \\ \forall b_1, b_2, b_3 : \mathcal{B}; g : \mathbb{N}_1 \bullet b_3 \sqsubset b_2 \sqsubset b_1 \Rightarrow \\ g \in \text{Granularity}(b_1, b_3) \Leftrightarrow \\ (\exists g_1, g_2 : \mathbb{N}_1 \bullet g = g_1 * g_2 \wedge \\ g_1 \in \text{Granularity}(b_1, b_2) \wedge \\ g_2 \in \text{Granularity}(b_2, b_3)) \end{array} \right.$$

Within a band, behaviour is defined using *events* (which are instantaneous), *activities* (that have duration) and *state* (both discrete and continuous). These are defined in later sections, but important here is the property that events are defined to be instantaneous. And two or more events may be defined to be simultaneous. A band's *precision* is the constraint on the duration of 'instantaneous' and 'simultaneous' when measured in a finer band. For example, if the precision of the hour band is defined to be one minute then simultaneous events must occur within a minute of each other.

As precision can only be expressed using the granularity of a finer band it is defined on a pair of (upper and lower) time bands.

$$\begin{array}{|l}
\hline
Precision : (\mathcal{B} \times \mathcal{B}) \leftrightarrow \mathbb{N}_1 \\
\hline
\text{dom}(Precision) = \{b1, b2 : \mathcal{B} \mid b2 \sqsubset b1\} \\
\forall b1, b2, b3 : \mathcal{B}; p : \mathbb{N}_1 \bullet b3 \sqsubset b2 \sqsubset b1 \Rightarrow \\
Precision(b1, b3) \leq Precision(b1, b2) * \min(Granularity(b2, b3))
\end{array}$$

The definition of precision enables the framework to be used effectively for requirements specification. A temporal requirement such as a deadline is band-specific; similarly the definition of a timing failure. For example, being one second late may be a crucial failure in a computing device, whereas on a human scale being one second late for a meeting is meaningless. The duration of an activity is also ‘imprecise’ (within the band). Stating that a job will take three months is assumed to mean plus or minus a couple of days. Of course the precision of a band can only be explored in a lower band.

Again with the lecturing example, assume the precision of the minute band is one second. The instantaneous ‘slide change’ event when mapped to a laptop activity in a lower band must have a duration of not more than one second.

A key aspect of the timeband framework is that certain entities are considered to be instantaneous, and that they are then mapped to actions that have duration in a more detailed description of the system. One means of achieving this property would be to give all such entities a distinct precision. However in constructing behaviours from collections of entities, composition is much more straightforward if the same notion of precision applies. Moreover, the property of being ‘instantaneous’ relates to the level of the temporal abstraction not to the event itself. Hence the timeband framework starts by defining the bands and then places entities into the bands. If the entity is instantaneous it is represented by an event; if it has duration then it is represented by an activity with a duration that is adequately expressed using the granularity of the chosen band. Hence ‘adequately’ means with sufficient (but not excessive) precision over the value of the defined duration.

As well as the system itself manifesting behaviour at many different time bands, the environment will exhibit dynamic behaviour at different granularities. The bands are therefore linked to the environment at the level determined by these dynamics. In many system abstractions it is useful to assume the environment is in some form of steady state. But this assumption is clearly false as environment evolves, perhaps as a result of the deployment of the embedded system under development. By mapping the rate of this evolutionary change to an appropriate (relatively slow) time band one can gain the advantage of the steady-state abstraction whilst not ignoring slower dynamics.

3.3 Classes of events/activities

In describing the behaviour of a system we often want to refer to repetitive activities/events. We say they are of a particular *class*, e.g., the event class corresponding a door opening, where the door will be opened many times during the life of the system. Each class of events/activities has a unique time band and name that we use to characterise the class.

\mathcal{C}
$band : \mathcal{B}$
$name : String$

The above is a Z schema: it defines a record type \mathcal{C} with two fields, $band$ and $name$.

3.4 Events

By definition, all actions within a band have similar dynamics. Within a band, *events* are instantaneous, while *activities* may have a non-zero duration. Events are a natural way of expressing change within a system. By first defining behaviours to be instantaneous, an abstract definition of their cause and effect can be given. Also seemingly impossible specifications can be given clear semantics. For example, the change-of-state event to turn of a water pump (as used in the case study in Section 7) is an event that ideally is instantaneous at some level of abstraction but clearly must take time at a more detailed level of description (in a finer band).

In a particular behaviour there may be any (countable) number of instances of events of a particular class, including zero. The set of instances of an event class within a behaviour are totally ordered by precedence (see below), and hence we can also assign a unique index to an event instance. An event instance, “event” for short, is characterised by its class (time band and name), and a natural number index, n , indicating that it is occurrence n of events of that class within a behaviour.

\mathcal{E}
$class : \mathcal{C}$
$index : \mathbb{N}$

We use the notation $c \# i$ to stand for the event of class c that has index number i . We define a shorthand for the time band of an event.

$band : \mathcal{E} \rightarrow \mathcal{B}$
$\forall e : \mathcal{E} \bullet band(e) = e.class.band$

For a band, b , $Events(b)$ defines the set of all events in that band.

$Events : \mathcal{B} \rightarrow \mathbb{P} \mathcal{E}$
$\forall b : \mathcal{B} \bullet Events(b) = \{e : \mathcal{E} \mid band(e) = b\}$

3.5 Precedence

Although time is of central important, there are contexts in which pure *order* is a more natural way of describing behaviour [16, 1, 11] (X was before Y, e.g., “before the end of the shift”, “after the plane took off”, “before the flood”, “after the thread

has completed”, “before the gate has fired”). The framework must therefore represent both precedence relations and temporal frames of reference.

There is of course a strong link between temporal order (i.e., time-stamped events and activities) and precedence relations. However, in this framework, we do not impose an equivalence between time and precedence. Due to issues of precision, time cannot be used to infer precedence unless the time interval between two events is sufficiently large in the band of interest.

Where bands are, at least partially, ordered by granularity, then order and hence potential causality is preserved as one moves from the finer to the coarser bands. However, order and hence causality are not necessarily maintained as one moves down through the bands. Where order is important then proof must be obtained by examining the inter-band relationships.

A precedence relation (\preceq) defines a partial ordering⁴ on the events. Only events in the same time band are related by the precedence relation. We use the operator \prec for strict precedence. A behaviour of a system will consist of some nonempty set of events ordered by precedence. The notation $(_ \preceq _)$ stands for the precedence relation taken as a whole.

<i>BehaviourEvents</i>
$ev : \mathbb{P}_1 \mathcal{E}$ $_ \preceq _ : \text{partial_order}[\mathcal{E}]$ $_ \prec _ : \mathcal{E} \leftrightarrow \mathcal{E}$
$(_ \preceq _) \in ev \leftrightarrow ev$ $\forall e, f : \mathcal{E} \bullet$ $(e \preceq f \Rightarrow \text{band}(e) = \text{band}(f)) \wedge$ $(e \prec f \Leftrightarrow e \preceq f \wedge e \neq f)$ $\forall c : \mathcal{C}; i, j : \mathbb{N} \bullet i < j \wedge c \# j \in ev \Rightarrow c \# i \in ev \wedge c \# i \prec c \# j$

The above Z schema declares a number of fields and constrains them to satisfy the predicate below the line. Note that we don't insist that all pairs of events are related one way or the other, but if both $e \preceq f$ and $f \preceq e$, because it is a partial order we insist that $e = f$. If there is an instance of a event with index j , then there must be instances for all indices less than j and these instances must precede the j th instance.

3.6 Simultaneous and immediate

In the specification of a system, an event may cause a response *immediately* (*instantaneously*) – meaning that at this band the response is within the precision of the band. This use of untimed notions helps eliminate the problem of over specifying requirements that is known to lead to implementation difficulties [13]. For example consider the naturally specified requirement ‘when the fridge door opens the light must come on immediately’; this apparently give no scope for an implementation to incorporate the necessary delays of switches, circuitry and the light's own latency. Making the

⁴ See Appendix A.

term ‘immediate’ band specific, enables a finer-granularity band to include the necessary delays, latencies and processing time that are needed to support the immediate behaviour at the higher band. This separation of concerns removes the need to add a precise deadline to the ‘light-on’ event. An explicit deadline (of say 8.5ms) is too concrete – rather the deadline is ‘the definition of immediate in this band’. Obviously being immediate in the hour band is not the same as being immediate in the microsecond band.

Two events may have a precedence relations (eg. slide X before slide Y) but occur at the same time (same hour).

It follows from these observations that, in this framework, there is a difference between two events being simultaneous and being ‘at the same time’. The former is a much stronger statement. Here two simultaneous events (in band B say) must, when viewed from a finer band, be within the precision of band B. Whereas ‘at the same time’ only required the two events to occur within the granularity of band B. As the precision is typically 1/10th to 1/100th of the granularity, clearly events being simultaneous is a much tighter constraint.

Precedence gives rise to potential causality. If P is before Q then information could flow between them, indeed P may be the cause of Q. In the use of the framework for specification we will need to use the stronger notion of precedence to imply causality. For example, “when the fridge door opens the light must come on”. Within the band of human experience this can be taken to be *immediate* (simultaneous but ordered). At a finer band a number of electro-mechanical activities will be needed to be described that will sense when the door is open and enable power to flow to the light. Importantly, no causality relationship can be inferred (without explicit precedence) for two events occurring at the same time within their particular band. In effect they are logically concurrent and may occur in sequence or overlapped in time when viewed from a lower band.

We introduce a separate relation (\simeq) to denote that two events are simultaneous. While \simeq is reflexive and symmetric, it isn’t transitive.⁵ One event, f , immediately follows another, e , written $e \preceq f$, if f both follows e and is simultaneous with e . Behaviours are extended to include simultaneous events. This schema includes schema *BehaviourEvents*, and hence includes all the fields of that schema as well as its constraints.

<i>BehaviourSimultaneous</i>
<i>BehaviourEvents</i>
$-\simeq- : \text{symmetric_rel}[\mathcal{E}]$
$-\preceq- : \mathcal{E} \leftrightarrow \mathcal{E}$
$(-\simeq-) \in \text{ev} \leftrightarrow \text{ev}$
$(\forall e, f, g, h : \mathcal{E} \bullet$
$(e \simeq f \Rightarrow \text{band}(e) = \text{band}(f)) \wedge$
$(e \preceq f \preceq h \wedge e \preceq g \preceq h \wedge e \simeq h \Rightarrow f \simeq g) \wedge$
$(e \preceq f \Leftrightarrow e \preceq f \wedge e \simeq f))$

⁵ See Appendix A for the definition of *symmetric_rel*.

3.7 Activities

An activity has a class (time band and name) and an instance number.

\mathcal{A}
<i>class</i> : \mathcal{C}
<i>instance</i> : \mathbb{N}

We also use the notation $c \# i$ to refer to the activity with class c and instance i . We define a shorthand for the time band of an activity.

$band : \mathcal{A} \rightarrow \mathcal{B}$
$\forall a : \mathcal{A} \bullet band(a) = a.class.band$

For a band, b , $Activities(b)$ defines the set of all activities in that band.

$Activities : \mathcal{B} \rightarrow \mathbb{P}\mathcal{A}$
$\forall b : \mathcal{B} \bullet Activities(b) = \{a : \mathcal{A} \mid band(a) = b\}$

An activity has associated with it a nonempty set of events (ie. $\mathbb{P}_1 \mathcal{E}$), all of which are in the same time band. Every activity, a , has associated with it a start event, $\uparrow a$, and possibly an end event, $\downarrow a$. An activity may not have an end event if it never terminates, or if we are only considering a partial trace of behaviour. For an activity of class c , the start events of such activities are of class $\uparrow c$ and the end events are of class $\downarrow c$; note that we have overloaded the up and down arrow symbols to function on both classes and activities. The start event of an activity should precede all events in the activity, which should themselves precede the activity's end event. The instance number of an activity is the same as the index number of its start event. Note that if we allow two activities of the same class to overlap, the instance number of an activity and the index number of its end event need not correspond.

<p><i>BehaviourActivities</i></p> <p><i>BehaviourSimultaneous</i></p> <p>$act : \mathbb{P} \mathcal{A}$</p> <p>$events_of : \mathcal{A} \mapsto \mathbb{P}_1 \mathcal{E}$</p> <p>$\uparrow : \mathcal{A} \mapsto \mathcal{E}$</p> <p>$\downarrow : \mathcal{A} \mapsto \mathcal{E}$</p> <hr/> <p>$events_of \in act \mapsto \mathbb{P}_1 ev$</p> <p>$dom(\uparrow) = act$</p> <p>$dom(\downarrow) \subseteq act$</p> <p>$\forall a : act \bullet$</p> <p>$events_of(a) \subseteq Events(band(a)) \wedge$</p> <p>$\uparrow a \in events_of(a) \wedge$</p> <p>$(a \in dom(\downarrow) \Rightarrow \downarrow a \in events_of(a)) \wedge$</p> <p>$(\uparrow a).class = \uparrow(a.class) \wedge$</p> <p>$(a \in dom(\downarrow) \Rightarrow (\downarrow a).class = \downarrow(a.class)) \wedge$</p> <p>$(\uparrow a).index = a.instance \wedge$</p> <p>$(\forall e : events_of(a) \bullet \uparrow a \preceq e \wedge (a \in dom(\downarrow) \Rightarrow e \preceq \downarrow a)) \wedge$</p> <p>$(\forall i : \mathbb{N} \bullet i < a.instance \Rightarrow a.class \# i \in act)$</p>

For the lecturing example, viewed at the year time band there may be an activity that corresponds to a course. The events of this activity include a set of lectures, all of which are after the start of the course and before its end. The lecture events may be related by precedence. At this level the precedence may just correspond to the dependence of material in one lecture on that in another, and hence the ordering of the lectures need not be total.

Many activities will have a repetitive cyclic behaviour with either a fixed periodicity or slowly varying pace. Other activities will be event-triggered. Most will have temporal constraints (deadlines). Activities are performed by *agents* (organizational, human or technical). In some bands all agents will be artificial (physical, computational or electrical), at others all human, and at others both will be evident. In addition to agents, there will often be the need for *resources* to enable the agent to make progress.

In this framework definition we will not include agents and resources; rather we concentrate on behaviour (events, activities and state). The scheduling of agents and resources so that activities meet their timing requirements is a natural extension to this description and would make use of standard scheduling and planning techniques.

3.8 Mappings between bands

In the components of the framework so far considered, all behaviours have been confined to a single band. In doing so some notions such as ‘instantaneous’, ‘simultaneous’, and ‘immediate’ have been defined but their semantic properties have not yet been fully defined. To do this, multiple-band behaviours need to be accommodated. This is achieved by mapping events in one band to activities in finer bands.




Fig. 1 Time Band Example

Events that are instantaneous in band A (say) may *map* to activities that have duration at some lower band with a finer granularity – we will denote this lower band as B. A key property of a band is the *precision* it defines for its time scale. This requires the activity associated with event E (in band A) to have a maximum duration of ρ (the precision of band A - as measured in band B). An illustration of a three band system with the mapping of events to activities is shown in Figure 1. As noted earlier, the start and end of an activity are themselves represented as events.

The link between any two bands is expressed in terms of each band's granularity and precision. Usually the finer of the two bands can be used to express these two measures for the broader band. Where physical time units are used for both bands these relations are straightforward. For example a band that is defined to have a granularity of an hour with a precision of two minutes is easily linked to a band with a granularity of ten seconds and precision of half a second. The granularity relation is a link from one time unit (1 hour) in the higher band to 360 units in the lower band. The precision of one minute means that a time reference at the higher band (e.g., 3 o'clock) will map down to the lower band to imply a time reference (interval) between 2.59 and 3.01. In general, two bands are said to be *ordered* if the precision of one band is larger than the granularity of the other.

If an event, e , maps to an activity, a , then that activity has a signature event, $sign(a)$, which corresponds to e at the lower level. Behaviours are extended with activities. The mapping preserves the precedence relation between two high-level events e and f by requiring that the signature events of the activity e is mapped to precedes the signature event of the activity f it is mapped to.

<p><i>BehaviourMapping</i></p> <p><i>BehaviourActivities</i></p> <p>$_ \rightsquigarrow _ : \mathcal{E} \leftrightarrow \mathcal{A}$</p> <p>$sign : \mathcal{A} \mapsto \mathcal{E}$</p> <hr/> <p>$(_ \rightsquigarrow _) \in ev \leftrightarrow act$</p> <p>$dom(sign) \subseteq act$</p> <p>$\forall e : ev; a : act \bullet e \rightsquigarrow a \Rightarrow (band(a) \sqsubset band(e) \wedge a \in dom(sign))$</p> <p>$\forall e : ev; a1, a2 : act \bullet band(a1) = band(a2) \wedge e \rightsquigarrow a1 \wedge e \rightsquigarrow a2 \Rightarrow a1 = a2$</p> <p>$\forall a : act \bullet a \in dom(sign) \Rightarrow sign(a) \in events_of(a)$</p> <p>$\forall e1, e2 : ev; a1, a2 : act \bullet$</p> <p style="padding-left: 2em;">$e1 \rightsquigarrow a1 \wedge e2 \rightsquigarrow a2 \wedge band(a1) = band(a2) \Rightarrow$</p> <p style="padding-left: 4em;">$(e1 \preceq e2 \Rightarrow sign(a1) \preceq sign(a2))$</p>

3.9 Durations

The function *duration* gives the time between any two events. To allow for lack of knowledge of the exact time between events and the granularity of the the time base, the result of *duration* is a time interval, i.e., a contiguous nonempty set of times, each of which is represented by a natural number.

$$Interval == \{I : \mathbb{P}_1 \mathbb{N} \mid (\forall t1, t2 : I; t : \mathbb{N} \bullet t1 < t < t2 \Rightarrow t \in I)\}$$

The duration of a terminating activity is determined from its start and end events.

<p><i>BehaviourDurations</i></p> <p><i>BehaviourMapping</i></p> <p>$duration : (\mathcal{E} \times \mathcal{E}) \mapsto Interval$</p> <p>$act_duration : \mathcal{A} \mapsto Interval$</p> <hr/> <p>$dom(duration) = \{e, f : ev \mid e \preceq f\}$</p> <p>$\forall e, f, g, h : ev \bullet e \preceq f \preceq g \preceq h \Rightarrow$</p> <p style="padding-left: 2em;">$(\forall I_1, I_2 : Interval \bullet duration(e, h) = I_1 \wedge duration(f, g) = I_2 \Rightarrow$</p> <p style="padding-left: 4em;">$I_2 \subseteq I_1)$</p> <p>$dom(act_duration) = dom(\downarrow)$</p> <p>$\forall a : dom(act_duration) \bullet act_duration(a) = duration(\uparrow a, \downarrow a)$</p>

Any event in the upper band is mapped to an activity in the lower band whose duration is within the precision of the upper band (with respect to the lower band).

Precision is not only important in defining the bounds on what it means for an event to be instantaneous (in a band), it is also used to constrain what is meant by two events to be simultaneous in some band. If e and f are simultaneous in band b (with precision ρ with respect to the lower band c) then the signature events of the mapped activities must occur within ρ in band c . Similarly, two events can be defined to be ‘not simultaneous’ and may require some component of the system to test that this erroneous situation does not occur. Again, by placing such a requirement in the

right band, the necessary tolerance on the implementation of the monitoring task is precisely defined.

<i>BehaviourPrecision</i>
<i>BehaviourDurations</i>
$\forall e : ev; a : act; n : \mathbb{N} \bullet$ $e \rightsquigarrow a \wedge n \in act_duration(a) \Rightarrow n \leq Precision(band(e), band(a))$ $\forall e1, e2 : ev; a1, a2 : act; n : \mathbb{N} \bullet e1 \simeq e2 \wedge$ $e1 \rightsquigarrow a1 \wedge e2 \rightsquigarrow a2 \wedge band(a1) = band(a2) \wedge$ $n \in duration(sign(a1), sign(a2)) \Rightarrow$ $n \leq Precision(band(e1), band(a1))$

3.10 Clocks

For the time bands associated with computational activity there is usually a strong notion of time and (adequately accurate) physical clocks that will aid scheduling and coordination. This is also increasingly the case with the bands of human experience as external sources of time and temporal triggers abound [17]. So measures such as second, minute, hour, day, week, month, year, decade and century are now universal. But other time scales such as ‘generation’, ‘era’ and ‘age’ are also used in specific domains. In a different context the granularity of a band may relate to a physical property of the system, such as the rotation of the crank shaft for an engine control unit.

A frame of reference defines an abstract clock that counts *ticks* of the band’s granularity and can be used to give a time stamp to events and activities. A band may have more than one such abstract clock but they progress at the same rate. For example the day band will have a different abstract clock in each distinct geographical time zone.

We develop a consistent model of time by representing certain moments in the dynamics of a band as “clock tick” events, which are modelled just like any other event. When necessary, an event can be situated in absolute time (within the context of a defined band and clock) by stating a precedence relationship between the event and one or more clock ticks. So an event occurred between 2.00 and 3.00 (in the hour band) if the event occurred after the start of hour from 2.00 to 3.00 but before the end of that hour. Note this is different to saying the event occurred ‘at 2.00’. Here the implication is that it is simultaneous with the 2.00 event. So ‘I will arrive at 2.00’ is satisfied by arriving sufficiently close to the 2.00 event (within the precision of the hour band). However ‘I’ll will arrive by 3.00’ is quite different and allows the arrival event to occur anytime up to the 3.00 event.

A clock can be modelled as a sequence of clock-tick events of a given class, and hence a given time band. Successive clock-tick events are separated by one time unit in the granularity of the band. They are therefore never simultaneous.

4 State

In modelling state within the timeband framework there are a number of issues we need to take into account:

- observations: within a particular time band only a subset of the state variables (observations) will be relevant,
- nondeterminism: at the time interval corresponding to an event within a band, there may be a set of possible values of a state variable,
- change-of-state events: for discrete state observations, changes in value correspond to change-of-state events,
- accuracy: for a continuous state observation there will be a maximum change over a time interval corresponding to the precision of the band, and
- rate-of-change: also for continuous observations there will be a maximum rate of change over an interval of size the granularity of the time band.

The observation variables of different time bands may be different. Typically, the representation becomes richer as one moves down to a time band with a finer granularity. We say that the state of the system is *projected* onto a band; in some bands not all possible observation variables will be accessible (as the time spent in that state is too short). To illustrate, consider an automatic door:

- at a high time band one can view the door as either *open* or *closed*, with “instantaneous” events to open or close it;
- at a lower time band the open and close events take time, and there are new activities *opening* and *closing*
- at a lower level still one may model how far open the door is by a percentage between 0% open (i.e., closed) and 100% open; this numeric measurement may either be discrete, with some granularity, or continuous; if it is discrete then, at a still lower time band, it may be discrete with a finer granularity.

This can be modelled by having different observation variables at different time bands.

4.1 States

The state space can be represented by a mapping from variables names taken from the set Var to values taken from the set Val .

$$State[Var, Val] == Var \rightarrow Val$$

The above definition of $State$ is generic in the set of variables and values. A state, $\sigma \in Var \rightarrow Val$, maps each variable name in its domain to its value in that state. For simplicity we use the universal set Val for all values, rather than each variable having values of a particular type.

The set of observation variables in a particular time band is fixed, and hence it is useful to refer to sets of states, all of which have the same variables (i.e., domain).

$$StateSet[Var, Val] == \{ss : \mathbb{P} State[Var, Val] \mid (\forall \sigma_1, \sigma_2 : ss \bullet \text{dom}(\sigma_1) = \text{dom}(\sigma_2))\}$$

For example, if x and y are variables and integers are values then s is a state and ss is a state set.

$$\begin{aligned} s &= \{x \mapsto 0, y \mapsto 1\} \\ ss &= \{\{x \mapsto 0, y \mapsto 1\}, \{x \mapsto 0, y \mapsto 0\}\} \end{aligned}$$

4.2 Nondeterministic state

Over the time interval corresponding to an event, e , within a particular time band, one can extract the set, ss , of actual states that occur in that interval. Unfortunately, the set of states view doesn't reflect the reality of observing multiple variables, all of which are evolving over time. For example, if we have two variables x and y which are both initially zero, and if in quick succession x changes to one and then y changes to one, then there is no point at which the state has x with value zero and y with value one. However, a program sampling the two variables may first sample x and get zero and then sample y and get one.

To address this issue we introduce the notion of a nondeterministic state, which for each variable records a set of values. This has less information than the equivalent set of states. The "sets of states" view is adequate for a single sequential process controlling all the variables in the state, but if there are concurrent processes or an externally evolving environment, observation of the state at a particular time precision may observe one variable at one instant and another at a slightly different instant. Hence, if we don't determine the order of observation of the variables, there is a set of values that we can observe for each variable at that time "instant". This leads to a less determined representation of the state, a *nondeterministic state*, in which each variable is mapped to a nonempty⁶ set of possible values.

$$NDState[Var, Val] == Var \mapsto \mathbb{P}_1 Val$$

Note that nondeterministic states are a form of state with values replaced by sets of values:

$$NDState[Var, Val] = State[Var, \mathbb{P}_1 Val]$$

As with states we define the sets of nondeterministic states, all of which have the same variables.

$$NDStateSet[Var, Val] == StateSet[Var, \mathbb{P}_1 Val]$$

The set of (standard) states that may be *apparent* in a nondeterministic state can be extracted by considering all possible states such that each variable maps to an element of its set of possible values.

$$\begin{array}{l} \boxed{\begin{array}{l} [Var, Val] \\ \text{apparent} : NDState[Var, Val] \rightarrow StateSet[Var, Val] \\ \forall ns : NDState[Var, Val] \bullet \\ \text{apparent}(ns) = \{\sigma : \text{dom}(ns) \rightarrow Val \mid (\forall v : \text{dom}(ns) \bullet \sigma(v) \in ns(v))\} \end{array}} \end{array}$$

⁶ Hence \mathbb{P}_1 rather than \mathbb{P} .

In the opposite direction, given a set of states, one may define a corresponding non-deterministic state.

$$\begin{array}{l}
 \hline \hline
 [Var, Val] \text{---} \\
 \hline
 nondet : StateSet[Var, Val] \rightarrow NDState[Var, Val] \\
 \hline
 \forall ss : StateSet[Var, Val] \bullet \\
 \quad \mathbf{let} \text{ vars} == \{v : Var \mid (\exists \sigma : ss \bullet v \in \text{dom}(\sigma))\} \bullet \\
 \quad \quad nondet(ss) = (\lambda v : \text{vars} \bullet \{y : Val \mid (\exists \sigma : ss \bullet y = \sigma(v))\})
 \end{array}$$

Note that while

$$ns = nondet(apparent(ns)) \quad (1)$$

in general, we only have

$$ss \subseteq apparent(nondet(ss)) \quad (2)$$

rather than equality. For example,

$$\begin{aligned}
 ss &= \{\{x \mapsto 0, y \mapsto 0\}, \{x \mapsto 1, y \mapsto 1\}\} \\
 ns &= nondet(ss) \\
 &= \{x \mapsto \{0, 1\}, y \mapsto \{0, 1\}\} \\
 ss' &= apparent(ns) \\
 &= \{\{x \mapsto 0, y \mapsto 0\}, \{x \mapsto 1, y \mapsto 1\}, \{x \mapsto 0, y \mapsto 1\}, \{x \mapsto 1, y \mapsto 0\}\} \\
 &\supset ss \\
 ns' &= nondet(ss') \\
 &= \{x \mapsto \{0, 1\}, y \mapsto \{0, 1\}\} \\
 &= ns
 \end{aligned}$$

Hence a set of states has potentially finer information than its corresponding non-deterministic state.

4.3 Behaviour with state

Within a behaviour one can ask for the nondeterministic state, $ev_val(e)$, that coincides with an event, e . For a discrete state variable, there is often a unique value, but if the event occurs close in time to a change of state then multiple values are possible to reflect our lack of knowledge of the actual value. For example, if an event is simultaneous with an hour band clock clk striking 12 then ev_val will return $\{clk \mapsto 11, clk \mapsto 12\}$. For a continuous state variable, if its value is changing at the time of the event then there is a range of values of the variable.

The set of nondeterministic states over an interval between (but not including) two events, $e1$ and $e2$, in the same time band is given by $interval_val(e1, e2)$. Behaviours are extended with state.

<p><i>BehaviourState</i></p> <p><i>BehaviourPrecision</i></p> <p>$observables : \mathcal{B} \rightarrow \mathbb{P} Var$</p> <p>$ev_val : \mathcal{E} \rightarrow NDState[Var, Val]$</p> <p>$interval_val : \mathcal{E} \times \mathcal{E} \rightarrow NDStateSet[Var, Val]$</p> <hr/> <p>$dom(ev_val) = ev$</p> <p>$dom(interval_val) = \{e1, e2 : ev \mid band(e1) = band(e2)\}$</p> <p>$\forall e : ev \bullet dom(ev_val(e)) = observables(band(e))$</p> <p>$\forall e1, e2 : ev \bullet (e1, e2) \in dom(interval_val) \Rightarrow$ $(\forall \sigma : interval_val(e1, e2) \bullet$ $dom(\sigma) = observables(band(e1)))$</p> <p>$\forall e1, e2, e3 : ev \bullet e1 \prec e2 \prec e3 \Rightarrow ev_val(e2) \in interval_val(e1, e3)$</p> <p>$\forall e1, e2 : ev \bullet e1 \preceq e2 \Rightarrow interval_val(e2, e1) = \{\}$</p> <p>$\forall e1, e2, e3, e4 : ev \bullet e1 \preceq e2 \preceq e3 \preceq e4 \Rightarrow$ $interval_val(e2, e3) \subseteq interval_val(e1, e4)$</p> <p>$\forall e1, e2 : ev \bullet e1 \prec e2 \Rightarrow$ $(\forall e : \{e1, e2\} \bullet$ $(\exists \sigma : interval_val(e1, e2) \bullet$ $(\forall v : dom(\sigma) \bullet ev_val(e)(v) \cap \sigma(v) \neq \{\})))$</p>
--

For a nonempty interval the states corresponding to the end-point events “overlap” in values with some state within the interval.

The set of nondeterministic states over an activity, a , is given by $act_val(a)$, which includes all the states between the start and end events of the activity, including at the start and end events.

<p><i>BehaviourStateActivities</i></p> <p><i>BehaviourState</i></p> <p>$act_val : \mathcal{A} \rightarrow NDStateSet[Var, Val]$</p> <hr/> <p>$dom(act_val) = act$</p> <p>$\forall a : act \bullet$ $(\forall \sigma : act_val(a) \bullet dom(\sigma) = observables(band(a))) \wedge$ $(\forall e1, e2 : events_of(a) \bullet$ $ev_val(e1) \in act_val(a) \wedge$ $interval_val(e1, e2) \subseteq act_val(a))$</p>
--

4.4 Change of state events

For discrete state, changes in value can be modelled by events. We can represent a change of state event in which variable v takes on the new value X , by the syntax $(v := X)$ for this class of event. A variable, v is constant between state change events for v .

<i>BehaviourStateChange</i>
<i>BehaviourStateActivities</i>
$\begin{aligned} \forall e_1 : ev; v : Var; X : Val \bullet e_1.class = (v := X) \Rightarrow \\ X \in ev_val(e_1)(v) \wedge \\ (\forall e_2 : ev \bullet e_1 \prec e_2 \wedge \\ \neg (\exists e : ev; Y : Val \bullet e_1 \prec e \prec e_2 \wedge e.class = (v := Y)) \Rightarrow \\ (\forall \sigma : interval_val(e_1, e_2) \bullet \sigma(v) = \{X\})) \end{aligned}$

4.5 Mapping states

If an event e is mapped to an activity a in a lower band, then the nondeterministic state at the event in the higher band corresponds to the union of the state values for the activity in the lower band.

<i>BehaviourStateMapping</i>
<i>BehaviourStateChange</i>
$\begin{aligned} \forall e : ev; a : act \bullet e \rightsquigarrow a \Rightarrow \\ (\forall v : observables(band(e)) \cap observables(band(a)) \bullet \\ ev_val(e)(v) = \{x : Val \mid (\exists ns : act_val(a) \bullet x \in ns(v))\}) \end{aligned}$

4.6 Accuracy and rates of change

As discussed above, within a single band a numeric-valued variable may have an accuracy and a maximum rate of change. Its accuracy is the maximum amount it can change over a period of size the precision of the band, and its maximum rate of change is the maximum amount it can change over a period of size the granularity of the band.

Within a particular time band, the rate of change of a state variable can be viewed as the change in its value over a time unit within the band. We'll illustrate this by discussing the maximum rate of change of a state variable, v .

- Within a given time band the maximum rate of change of v may be r , i.e., v can change by at most r over one time unit in that band.
- For a lower time band to be consistent with the upper band, the sum of the changes over a sequence of time units within the lower time band, with length corresponding to the granularity of the upper band with respect to the lower band, must be at most r .

<p><i>BehaviourRates</i></p> <hr/> <p><i>BehaviourStateMapping</i></p> <p>$accuracy : \mathcal{B} \rightarrow (Var \leftrightarrow \mathbb{R})$</p> <p>$rate : \mathcal{B} \rightarrow (Var \leftrightarrow \mathbb{R})$</p> <hr/> <p>$\forall b : \mathcal{B} \bullet$</p> <p>$\text{dom}(accuracy(b)) = \text{dom}(rate(b)) \wedge \text{dom}(rate(b)) \subseteq observables(b)$</p> <p>$\forall e : ev \bullet (\forall v : \text{dom}(accuracy(band(e))) \bullet$</p> <p>$(\forall x, y : ev_val(e)(v) \bullet abs(x - y) \leq accuracy(band(e))(v)))$</p> <p>$\forall e1, e2 : ev \bullet e1 \prec e2 \Rightarrow (\forall v : \text{dom}(rate(band(e1))) \bullet$</p> <p>$(\forall \sigma_1, \sigma_2 : interval_val(e1, e2) \bullet (\forall x : \sigma_1(v); y : \sigma_2(v) \bullet$</p> <p>$abs(x - y) \leq rate(band(e1))(v) * max(duration(e1, e2))))))$</p> <p>$\forall e : ev; a : act \bullet e \rightsquigarrow a \Rightarrow$</p> <p>$(\forall v : \text{dom}(accuracy(band(e))) \cap \text{dom}(rate(band(a))) \bullet$</p> <p>$rate(band(a))(v) * max(act_duration(a)) \leq accuracy(band(e))(v))$</p>
--

Satisfying the consistency condition has a special case if we consider the state variable to be uniform between two levels, or *uniform* for short. If the granularity of the upper time band with respect to the lower time band is n , then for a uniform state variable the maximum rate of change in the lower time band will be at most r/n . At a still lower time band with granularity m with respect to the above lower band (and hence granularity $m * n$ with respect to the upper time band) the maximum rate of change in this still lower time band will be at most $r/(m * n)$. With a uniform state variable, as the size of the time unit of the band approaches zero the rate of change approaches the derivative of the state variable with respect to time.

5 Summary

Rather than have a single notion of time, the proposed framework allows a number of distinct time bands to be used in the specification or description of a system. System behaviours are always relative to (defined within) a band.

The above discussion has defined the timeband framework and introduced a number of key notions that are central to the framework. Here we summarize these ideas:

- *band* – a subset of system behaviours (discrete and continuous) with similar temporal properties;
- *system* – a partially ordered set of bands;
- *separation* – the property of being able to assume that activities in lower (quicker) bands are instantaneous and the state of higher (slower) bands is unchanging;
- *granularity* – the unit of time defined by a band;
- *precision* – the constraint on instantaneous behaviour within a band;
- *event* – an instantaneous action within a band;
- *activity* – an action with duration within a band;
- *duration* – a time interval between events;
- *clock* – an abstract band-specific clock that produces ticks (events) at the granularity of the band;

- *precedence* – one event happening after another event;
- *simultaneous* – two events occurring at the same instant;
- *immediate* – a precedence relation between two simultaneous events;
- *mapping* – a link between an event in one band to an activity in a lower band;
- *state* – the observations available within a time band;
- *nondeterministic state* – the observations over the period of an event;
- *change-of-state events* – for discrete observations;
- *accuracy* – maximum “instantaneous” change in a continuous variable;
- *rate-of-change* – maximum rate of change of a continuous variable over a unit of time.

6 Towards a Language for Timebands

Having presented a model for the timeband framework it is then necessary to define a language that can be used to specify system requirements and behaviour. Such a language is derived from the abstract model. In this paper we do not attempt to provide a single, or even a complete, timeband language. Rather we illustrate features that such a language could usefully contain. These are used in a short example of the use of time bands in the Section 7.

6.1 Predicates

We represent predicates over a state space, Σ , via subsets of Σ .

Definition 1 (Predicate) For a state space Σ , a predicate is represented by a set of states.

$$Pred[\Sigma] == \mathbb{P} \Sigma$$

We use the conventional notations, “ \wedge ”, “ \vee ”, and “ \neg ” instead of intersection, union, and complement of sets, respectively, when dealing with predicates.

6.2 Predicates on sets of states

Given a predicate, p , there are two obvious ways to promote it to a set of states (as in modal logics [12]). If p holds for *all states* in the set, we write $\boxtimes p$, and if p holds for *some state* in the set, we write $\boxdot p$.

Definition 2 (All states and some states)

$\begin{aligned} \boxtimes &: Pred[\Sigma] \rightarrow Pred[\mathbb{P}_1 \Sigma] \\ \boxdot &: Pred[\Sigma] \rightarrow Pred[\mathbb{P}_1 \Sigma] \end{aligned}$
$\begin{aligned} \forall p : Pred[\Sigma]; ss : \mathbb{P}_1 \Sigma \bullet \\ ss \in (\boxtimes p) &\Leftrightarrow (\forall \sigma : ss \bullet \sigma \in p) \wedge \\ ss \in (\boxdot p) &\Leftrightarrow (\exists \sigma : ss \bullet \sigma \in p) \end{aligned}$

We promote the boolean operators to predicates on sets of states in the obvious way (because they are defined as predicates, but over sets of states rather than states). We have the following properties.

$$\begin{aligned}
\neg \boxtimes p &= \square(\neg p) \\
\boxtimes p &\Rightarrow \square p \\
\boxtimes p \wedge \boxtimes q &= \boxtimes(p \wedge q) \\
\square p \vee \square q &= \square(p \vee q) \\
\boxtimes p \vee \boxtimes q &\Rightarrow \boxtimes(p \vee q) \\
\square(p \wedge q) &\Rightarrow \square p \wedge \square q
\end{aligned}$$

Note that the second property above is valid because the sets of states must be non-empty.

From these properties we can deduce other properties such as the following.

$$(\square p \Rightarrow \boxtimes q) \Rightarrow (\boxtimes(p \Rightarrow q)) \quad (3)$$

This can be shown as follows.

$$\begin{aligned}
&\square p \Rightarrow \boxtimes q \\
&\equiv \neg \square p \vee \boxtimes q \\
&\equiv \boxtimes \neg p \vee \boxtimes q \\
&\Rightarrow \boxtimes(\neg p \vee q) \\
&\equiv \boxtimes(p \Rightarrow q)
\end{aligned}$$

6.3 Predicates on nondeterministic states

A nondeterministic predicate is a predicate on a nondeterministic state.

Definition 3 (Nondeterministic predicate)

$$NDPred[Var, Val] == Pred[NDState[Var, Val]]$$

We promote a predicate, p , on a single state, to a predicate on a nondeterministic state in two ways. If p holds for all apparent states (see Section 4.2) derivable from the nondeterministic state, we say p *definitely* holds for the nondeterministic state, written $\boxtimes p$, and if p holds for at least one apparent state derivable from the nondeterministic state, we say p *possibly* holds for the nondeterministic state, written $\odot p$.

Definition 4 (Definitely and possibly)

$ \begin{aligned} &\boxtimes : Pred[State[Var, Val]] \rightarrow NDPred[Var, Val] \\ &\odot : Pred[State[Var, Val]] \rightarrow NDPred[Var, Val] \end{aligned} $
$ \begin{aligned} &\forall p : Pred[State[Var, Val]]; ns : NDState[Var, Val] \bullet \\ &\quad ns \in (\boxtimes p) \Leftrightarrow (\forall \sigma : apparent(ns) \bullet \sigma \in p) \wedge \\ &\quad ns \in (\odot p) \Leftrightarrow (\exists \sigma : apparent(ns) \bullet \sigma \in p) \end{aligned} $

We promote the boolean operators to nondeterministic predicates in the obvious way (because nondeterministic predicates are defined as predicates, but over nondeterministic states rather than states).

For example, in the methane control example considered in the following section, if the methane in a mine shaft is ever over a critical level, then to avoid causing an explosion, the pump extracting water from the mine must be off. We can formalise this property by the following nondeterministic predicate, that characterises those nondeterministic states which satisfy the property.

$$\odot(\text{methane} \geq \text{Critical}) \Rightarrow \otimes(\text{pump} = \text{Off})$$

If the methane is critical at some instant (i.e., the nondeterministic state includes an apparent state in which the methane is critical), then the pump is definitely off (i.e., it is off for all apparent states of the nondeterministic state).

6.4 Properties of nondeterministic predicates

The definitely and possibly operators are related to “all states” and “some state” as follows.

$$ns \in \otimes p \Leftrightarrow \text{apparent}(ns) \in \boxtimes p \quad (4)$$

$$ns \in \odot p \Leftrightarrow \text{apparent}(ns) \in \boxminus p \quad (5)$$

Hence, we have the following properties directly derivable from the properties of predicates on sets of states.

$$\neg \otimes p = \odot(\neg p) \quad (6)$$

$$\otimes p \Rightarrow \odot p \quad (7)$$

$$\otimes p \wedge \otimes q = \otimes(p \wedge q) \quad (8)$$

$$\odot p \vee \odot q = \odot(p \vee q) \quad (9)$$

$$\otimes p \vee \otimes q \Rightarrow \otimes(p \vee q) \quad (10)$$

$$\odot(p \wedge q) \Rightarrow \odot p \wedge \odot q \quad (11)$$

There are two interesting properties of definitely and possibly that don't hold for “all states” and “some states”.

Theorem 1 *If the free variables of p and q are disjoint, then*

$$\odot p \wedge \odot q = \odot(p \wedge q)$$

$$\otimes(p \vee q) = \otimes p \vee \otimes q$$

Proof. We focus on the first property because the second can be derived from it because $\otimes p = \neg \odot \neg p$. The reverse implication is property (11) above. In the forward direction, if we assume $ns \in \odot p$ and $ns \in \odot q$, then $(\exists \sigma : \text{apparent}(ns) \bullet \sigma \in p)$ and $(\exists \sigma : \text{apparent}(ns) \bullet \sigma \in q)$. Hence let $\sigma_1 \in p \cap \text{apparent}(ns)$, $\sigma_2 \in q \cap \text{apparent}(ns)$, w be the set of free variables occurring in p , and $\sigma = (w \triangleleft \sigma_1) \cup (w \triangleleft \sigma_2)$, where $w \triangleleft \sigma$ is the state σ restricted to just the observations in w and $w \triangleleft \sigma$ is σ restricted to the observations not in w . Because p depends only on variables in w and $w \triangleleft \sigma = w \triangleleft \sigma_1$, it follows that $\sigma \in p$. Similarly, because q only depends on variables not in w , $\sigma \in q$. Finally, because both σ_1 and σ_2 are in $\text{apparent}(ns)$, $\sigma \in \text{apparent}(ns)$. Hence $(\exists \sigma : \text{apparent}(ns) \bullet \sigma \in p \wedge q)$, i.e., $ns \in \odot(p \wedge q)$.

□

Note that these properties hold for $\odot p$ but not $\Box p$. For example, if

$$ss = \{\{x \mapsto 0, y \mapsto 0\}, \{x \mapsto 1, y \mapsto 1\}\}$$

we have $ss \in \Box(x = 0) \wedge \Box(y = 1)$ but not $ss \in \Box(x = 0 \wedge y = 1)$.

6.5 On the relation between reality and observation

Over the time interval corresponding to an event, e , within a particular time band, one can extract the set, ss , of actual states that occur in that interval. The corresponding nondeterministic state is $\text{nondet}(ss)$. Note from property (2) in Section 4.2 that $ss \subseteq \text{apparent}(\text{nondet}(ss))$. If we assume $\text{nondet}(ss) \in \otimes p$, then we have by (4) that $\text{apparent}(\text{nondet}(ss)) \in \otimes p$ and hence $ss \in \otimes p$, i.e., p holds for all actual states corresponding to event e .

If the environment satisfies p at some point in the interval, i.e., $ss \in \Box p$, then as $ss \subseteq \text{apparent}(\text{nondet}(ss))$ we have $(\exists \sigma : \text{apparent}(\text{nondet}(ss)) \bullet \sigma \in p)$, and hence $\text{nondet}(ss) \in \odot p$. That is, if reality satisfies $\Box p$, then our (weaker) nondeterministic view satisfies $\odot p$.

In the simple case in which we are only dealing with one free variable, x , in a predicate, e.g., the predicate is of the form $\otimes(x \in S)$ or $\odot(x \in S)$, where S is constant over the observation interval, then $\otimes(x \in S)$ in the nondeterministic view is equivalent to $\otimes(x \in S)$ in reality, and $\odot(x \in S)$ in the nondeterministic view is equivalent to $\Box(x \in S)$ in reality. Special cases of these predicates are comparisons of a variable with an expression, C , that is constant over the observation interval, e.g., $x = C$ and $x < C$.

When one samples a variable, x , in the environment, and gets the value C , one can deduce $\Box(x = C)$, which is equivalent to $\odot(x = C)$. Similarly, by sampling y we may deduce $\Box(y = D)$, which is equivalent to $\odot(y = D)$. These samples allow one to deduce $\odot(x = C \wedge y = D)$ but not the stronger condition $\Box(x = C \wedge y = D)$. This formalises the property that sampling two boolean variables, *top* and *bottom*, introduced in Section 2. Getting two sample values, e.g., *true* and *true*, does not allow one to deduce that the two variables simultaneously have those values (i.e., we can't

deduce $\Box(top \wedge bottom)$, but we can deduce the weaker property $\odot(top \wedge bottom)$. We have

$$ss \in \Box(top \wedge bottom) \Rightarrow nondet(ss) \in \odot(top \wedge bottom)$$

but not the other way around, in general. Note that $\odot(top \wedge bottom) = \odot top \wedge \odot bottom$, but only $\Box(top \wedge bottom) \Rightarrow \Box top \wedge \Box bottom$, in general.

6.6 Application to state model

In this section we develop some notation for using nondeterministic predicates with the nondeterministic state model. Given a behaviour and an event e , $ev_val(e)$ gives the nondeterministic state corresponding to event e . For a nondeterministic predicate, p , we introduce the notation $p @ e$ to state that p holds for the nondeterministic state corresponding to e , i.e.,

$$p @ e \Leftrightarrow ev_val(e) \in p$$

For events $e1$ and $e2$, $interval_val(e1, e2)$ gives the set of nondeterministic states occurring (strictly) between the two events. We introduce the notation p during $(e1, e2)$ to stand for p holding for all nondeterministic states between $e1$ and $e2$, and p within $(e1, e2)$ to state that p holds for some nondeterministic state between $e1$ and $e2$, i.e.,

$$\begin{aligned} p \text{ during } (e1, e2) &\Leftrightarrow interval_val(e1, e2) \in \boxtimes p \wedge \\ p \text{ within } (e1, e2) &\Leftrightarrow interval_val(e1, e2) \in \Box p \end{aligned}$$

The properties of behaviours on state allow one to deduce properties expressed in terms of these relations. For example, if a property definitely holds at an end point event of a nonempty interval then, because the state of the end point “overlaps” with those in the interval, we get the following property for all pairs of events $e1$ and $e2$ such that $e1 \prec e2$,

$$(\boxtimes p) @ e1 \Rightarrow (\odot p) \text{ within } (e1, e2) \quad (12)$$

For example, if $\boxtimes(m \geq C) @ e1$ holds, i.e., for all values of m in the nondeterministic state corresponding to $e1$ we have $m \geq C$ holding, then because the end of event $e1$ corresponds to the start of the interval $(e1, e2)$, we have that $m \geq C$ at the very start of the interval $(e1, e2)$, but note that we don’t have any guarantee that m stays above C for any given period—not even the precision of the band—during the interval, and hence we can only deduce $\odot(m \geq C)$ holds within the interval.

We introduce two further shorthands to state that a nondeterministic predicate holds just before (after) an event.

$$\begin{aligned} p \text{ before } e &\Leftrightarrow (\exists e' : ev \bullet e' \preceq e \wedge p \text{ during } (e', e)) \wedge \\ p \text{ after } e &\Leftrightarrow (\exists e' : ev \bullet e \preceq e' \wedge p \text{ during } (e, e')) \end{aligned}$$

We assume all these relations have higher precedence than logical operators, but lower precedence than the other operators.

From the properties for predicates on sets of states we can deduce properties for *during* and *within*, e.g., the following property.

$$\begin{aligned} (p \text{ within } (e1, e2)) \Rightarrow (q \text{ during } (e1, e2)) &\Rightarrow \\ (p \Rightarrow q) \text{ during } (e1, e2) &\end{aligned} \quad (13)$$

This holds because

$$\begin{aligned} &(p \text{ within } (e1, e2)) \Rightarrow (q \text{ during } (e1, e2)) \\ \equiv &\text{ letting } sns = \text{interval_val}(e1, e2) \\ &sns \in \Box p \Rightarrow sns \in \boxtimes q \\ \equiv & \\ &sns \in (\Box p \Rightarrow \boxtimes q) \\ \Rightarrow &\text{ by (3)} \\ &sns \in \boxtimes(p \Rightarrow q) \\ \equiv & \\ &(p \Rightarrow q) \text{ during } (e1, e2) \end{aligned}$$

7 Example: Mine pump

The mine pump case study has been used by a number of formal frameworks to investigate and illustrate different specification approaches (see [15] for a number of examples). In its briefest form the case study involves two subsystems: a methane monitoring subsystem that sounds an alarm if the sensed level of methane is above a threshold, and a pump control subsystem that pumps water from the mine sump if the water level reaches a *High_water* level sensor (the pump then operates until a *Low_water* level sensor is reached). The two subsystems are coupled by the safety requirement not to operate the pump if the methane is high (due to risk of gas explosion). There is also a performance requirement that limits the number of days lost due to flooding to be two or less. In the following partial treatment we concentrate on the methane control subsystem.

We specify the mine pump in terms of *rely* and *guarantee* conditions [14], which are similar to preconditions and postconditions, except that *rely* and *guarantee* conditions are specified over the interval during which the system is running, rather than in terms of the before and after states for pre-/post-conditions. A *guarantee* condition is a condition that the system should maintain over its operating interval, provided the *rely* conditions hold for that interval.

Mine pump guarantee. A pump is used to extract water from a mine shaft. However, if there is a critical level of methane in the mine shaft, an explosion could result if the mine pump is operated. Hence, one requirement is that at all times while the system is running, it should guarantee that the pump is off at any time the methane level is critical.⁷

⁷ In a more detailed analysis, one may want to distinguish between the pump being turned off and it actually having come to a stop.

For a guarantee about a system's behaviour, if one can show the system guarantees $\circledast p$ holds for the nondeterministic view of the state, then $\boxtimes p$ holds for the corresponding real states. Hence in the mine pump example, one guarantee is

$$\boxtimes(\text{methane} \geq \text{Critical} \Rightarrow \text{pump} = \text{Off})$$

should hold for all (normal) states of the system, which is implied by the following requirement on all nondeterministic states of the system

$$\circledast(\text{methane} \geq \text{Critical} \Rightarrow \text{pump} = \text{Off})$$

which is equivalent to

$$\circledcirc(\text{methane} \geq \text{Critical}) \Rightarrow \circledast(\text{pump} = \text{Off}) \quad (14)$$

by the following reasoning.

$$\begin{aligned} & \circledast(\text{methane} < \text{Critical} \vee \text{pump} = \text{Off}) \\ \equiv & \text{Theorem 1; free variables in the two disjuncts are disjoint} \\ & \circledast(\text{methane} < \text{Critical}) \vee \circledast(\text{pump} = \text{Off}) \\ \equiv & \neg \circledcirc(\text{methane} \geq \text{Critical}) \vee \circledast(\text{pump} = \text{Off}) \\ \equiv & \circledcirc(\text{methane} \geq \text{Critical}) \Rightarrow \circledast(\text{pump} = \text{Off}) \end{aligned}$$

The “definitely” operator is interpreted with respect to a particular time band. Interestingly, from the point of view of this guarantee, all we require is that there is some time band in which this is satisfied (and this can be determined by the implementation). If we added a requirement ensuring that the pump is on whenever the methane level is safe (i.e., it is below critical by some bound), then this would constrain the choice of implementation time band.

Mine pump rely. In order to implement the requirement, one may rely on a number of properties of the environment. We'll assume that the implementation will be operating within a particular time band, but the properties we rely on are valid over a range of possible choices of time bands. Assume the accuracy of the methane level within the band is *Acc_meth*. Recall that this is the maximum amount the methane level can vary over a period of size the precision of the time band. Hence for any nondeterministic state of the system and any value Z

$$\circledcirc(\text{methane} = Z) \Rightarrow \circledast(\text{methane} \in Z \pm \text{Acc_meth}) \quad (15)$$

from which we can deduce

$$\circledcirc(\text{methane} \geq Z) \Rightarrow \circledast(\text{methane} \geq Z - \text{Acc_meth}) \quad (16)$$

Note that this accuracy is only concerned with the timing precision of the band. We should separately consider the accuracy of the methane sampling sensor itself, but for the purposes of this example we assume there is no sampling error.

Assume the maximum rate of change of methane level in the band is *Rate_meth*. Recall that this is the maximum change in the methane over a period of size the granularity of the time band. For an interval of size n , the level of methane can change

by at most n times $Rate_meth$ over that interval. Hence, for some value, Z , if the methane is ever at least than $Z + n * Rate_meth$ within the interval, then it must have been at least Z at the start of the interval. For all pairs of events e_1 and e_2 within the time band such that $e_1 \prec e_2$ and $max(duration(e_1, e_2)) \leq n$, and for any value Z ,

$$\begin{aligned} \odot(methane \geq Z + n * Rate_meth) \text{ within } (e_1, e_2) &\Rightarrow \\ \odot(methane \geq Z) @ e_1 & \end{aligned} \quad (17)$$

Mine pump implementation. The implementation samples the methane level at regular intervals. The i th sampling event is $s \# i$. It is assumed that the pump is off before the first sample.

$$\otimes(pump = Off) \text{ before } s \# 0 \quad (18)$$

We require that the maximum time between samples is n time units in the granularity of the implementation band, i.e., $max(duration(s \# i, s \# i + 1)) \leq n$. A sample event, $s \# i$, corresponds to an activity, $sample \# i$, at a lower time band. This activity, which all takes place within the precision of the upper time band, contains events to set up analog-to-digital converters to read the methane and water levels, wait until the conversions are complete and read the levels, and turn the pump off if the methane is above a threshold value. In addition, if the methane is below the threshold it turns the pump on or off depending upon the current water level.

Here we focus on the properties relating to the methane monitoring. If the pump is on up until sample i then,

- if the methane is definitely above the threshold for sample i , the *pump* must be off immediately after the sample

$$\otimes(methane \geq Threshold) @ s \# i \Rightarrow \otimes(pump = Off) \text{ after } s \# i \quad (19)$$

- if the *pump* is off immediately after sample i , the methane was possibly over the threshold during the sample,

$$\otimes(pump = Off) \text{ after } s \# i \Rightarrow \odot(methane \geq Threshold) @ s \# i \quad (20)$$

Note this is an example of the “must/can” sampling issue raised in Section 2.

Between any two successive sampling events the state of the pump is stable. Hence if the pump is off after sample i it remains off until at least sample $i + 1$:

$$\otimes(pump = Off) \text{ after } s \# i \Rightarrow \otimes(pump = Off) \text{ during } (s \# i, s \# i + 1) \quad (21)$$

From (19) and (21) one can deduce

$$\otimes(methane \geq Threshold) @ s \# i \Rightarrow \otimes(pump = Off) \text{ during } (s \# i, s \# i + 1) \quad (22)$$

If the pump was already off up until sample i then at sample i , if the methane is definitely at least the threshold, the pump remains off for the sample.

$$\begin{aligned} \otimes(pump = Off) \text{ before } s \# i \wedge \otimes(methane \geq Threshold) @ s \# i \\ \Rightarrow \otimes(pump = Off) @ s \# i \end{aligned} \quad (23)$$

To allow for the maximum rate of change between the sampling events (which are at most n time units apart) and the inaccuracies of the sampling events at each end of the interval, we require the following constraint between *Critical* and *Threshold*.

$$Critical \geq Threshold + n * Rate_meth + 2 * Acc_meth \quad (24)$$

Satisfaction of the specification.

Theorem 2 *The implementation fulfils the guarantee (14) whenever the system is running, provided the rely conditions hold.*

Proof. We assume the rely conditions and the conditions specified for the implementation and show that the guarantee condition (14) holds over the operating interval, by showing that it holds for every interval between one sample and the next, and at every sampling event. That is, for all natural numbers i

$$\odot(\text{methane} \geq \text{Critical}) \Rightarrow \otimes(\text{pump} = \text{Off}) \text{ @ } s \# i \quad (25)$$

$$\odot(\text{methane} \geq \text{Critical}) \Rightarrow \otimes(\text{pump} = \text{Off}) \text{ during } (s \# i, s \# i + 1) \quad (26)$$

First we show

$$\begin{aligned} \odot(\text{methane} \geq \text{Critical} - \text{Acc_meth}) \text{ within } (s \# i, s \# i + 1) &\Rightarrow \\ \otimes(\text{pump} = \text{Off}) \text{ during } (s \# i, s \# i + 1) &\quad (27) \end{aligned}$$

as follows

$$\begin{aligned} &\odot(\text{methane} \geq \text{Critical} - \text{Acc_meth}) \text{ within } (s \# i, s \# i + 1) \\ \Rightarrow &\text{from the constraint on the threshold (24)} \\ &\odot(\text{methane} \geq \text{Threshold} + n * \text{Rate_meth} + \text{Acc_meth}) \text{ within } (s \# i, s \# i + 1) \\ \Rightarrow &\text{from the maximum rate-of-change of methane (17)} \\ &\odot(\text{methane} \geq \text{Threshold} + \text{Acc_meth}) \text{ @ } s \# i \\ \Rightarrow &\text{from methane accuracy (16)} \\ &\otimes(\text{methane} \geq \text{Threshold}) \text{ @ } s \# i \\ \Rightarrow &\text{as the pump is turned off if methane is high (22)} \\ &\otimes(\text{pump} = \text{Off}) \text{ during } (s \# i, s \# i + 1) \end{aligned}$$

To show (25) we first consider the case for sampling events other than sample zero.

$$\begin{aligned} &\odot(\text{methane} \geq \text{Critical}) \text{ @ } s \# i + 1 \\ \Rightarrow &\text{from methane accuracy (16)} \\ &\otimes(\text{methane} \geq \text{Critical} - \text{Acc_meth}) \text{ @ } s \# i + 1 \\ \Rightarrow &\text{by (12)} \\ &\odot(\text{methane} \geq \text{Critical} - \text{Acc_meth}) \text{ within } (s \# i, s \# i + 1) \\ \Rightarrow &\text{from above lemma (27)} \\ &\otimes(\text{pump} = \text{Off}) \text{ during } (s \# i, s \# i + 1) \\ \Rightarrow &\text{by the definition of before} \\ &\otimes(\text{pump} = \text{Off}) \text{ before } s \# i + 1 \\ \Rightarrow &\text{by (23) and } \text{Critical} \geq \text{Threshold} + \text{Acc_meth} \\ &\otimes(\text{pump} = \text{Off}) \text{ @ } s \# i + 1 \end{aligned}$$

For sample zero we have $\otimes(\text{pump} = \text{Off}) \text{ before } s \# 0$ from (18), and hence it is sufficient to apply the equivalent of the last step of the above proof.

To show (26), we first note that, using (13), it is implied by

$$\begin{aligned} \odot(\text{methane} \geq \text{Critical}) \text{ within } (s \# i, s \# i + 1) &\Rightarrow \\ \otimes(\text{pump} = \text{Off}) \text{ during } (s \# i, s \# i + 1) &\end{aligned}$$

which we show as follows.

$$\begin{aligned}
& \odot(\text{methane} \geq \text{Critical}) \text{ within } (s \# i, s \# i + 1) \\
\Rightarrow & \text{ from methane accuracy (16)} \\
& \otimes(\text{methane} \geq \text{Critical} - \text{Acc_meth}) \text{ within } (s \# i, s \# i + 1) \\
\Rightarrow & \text{ as } \otimes p \Rightarrow \odot p \\
& \odot(\text{methane} \geq \text{Critical} - \text{Acc_meth}) \text{ within } (s \# i, s \# i + 1) \\
\Rightarrow & \text{ using lemma (27)} \\
& \otimes(\text{pump} = \text{Off}) \text{ during } (s \# i, s \# i + 1)
\end{aligned}$$

□

The above reasoning has been generic with respect to the choice of time band, in particular, the precision and granularity of the time band and the choice of the sampling interval n . From (24) we have

$$\text{Critical} - \text{Threshold} \geq n * \text{Rate_meth} + 2 * \text{Acc_meth}.$$

The larger the gap between *Critical* and *Threshold* the less time the pump will be active when it is safe to be active. The predominant factor is the choice of the sampling interval, here represented by n times the granularity of the band. If we add a further requirement that the pump should be active whenever the methane level is below a level of *Safe*, then this will put an upper bound on $n * \text{Rate_meth} + 2 * \text{Acc_meth}$.

This mine pump example does not need time bands in order to specify its behaviour. But the brief outline here does have a much simpler form than other descriptions. This clarity comes from being able to initially specify all behaviours as instantaneous events and all actions as immediate. This very real-time system did not need time in its initial specification, but it does allow time to play its full role in more detailed levels of description.

8 Conclusion

In this paper we have argued that complex real-time systems exhibit behaviour at many different time levels and that a useful aid in describing and specifying such behaviour is to use time bands. Viewing a system as a collection of event and activities within a finite set of bands is an effective means of separating concerns and identifying inconsistencies between different ‘layers’ of the system. Time bands are not mapped on to a single notion of physical time. Within a system there will always be a relation between bands but the bands need not be tightly synchronised. There is always some level of imprecision between any two adjacent bands.

The use of the timeband framework is intended to help develop a comprehensive foundation to the study and development of future systems. Of course an adequately expressive model of time is just one element of such a foundation, but it is perhaps the most important to define if dependable systems are to be engineered.

Acknowledgements The authors would like to thank the useful contributions of all the people who have discussed time bands with us – in particular Colin Fidge, Cliff Jones, Kun Wei and Jim Woodcock. This work is supported, in part, by the EPSRC(UK) through its funding of the INDEED project and the TrAmS platform, and is also supported, in part, by the Australian Research Council (ARC) Discovery Grant DP0987452, *Combining Time Bands and Teleo-Reactive Programs for Advanced Dependable Real-Time Systems*.

References

1. J. Allen. Towards a general theory of actions and time. *Artificial Intelligence*, 23:123–154, 1984.
2. G. Baxter, A. Burns, and K. Tan. Evaluating timebands as a tool for structuring the design of socio-technical systems. In P. Bust, editor, *Contemporary Ergonomics 2007*, pages 55–60. Taylor and Francis, 2007.
3. M. Bergadaà. Temporal frameworks and individual cultural activities: Four typical profiles. *Time and Society*, 16(2/3):387–408, 2007.
4. S. Blount, M.J. Walker, and S. Leroy. Coping with temporal uncertainty: when rigid ambitious deadlines don't make sense. In *Organization at the Limit*, pages 122–34. Springer, 2007.
5. A. Burns and G.D. Baxter. Time bands in systems structure. In *Structure for Dependability*, pages 74–90. Springer, 2006.
6. A. Burns, I.J. Hayes, G.Baxter, and C.J. Fidge. Modelling temporal behaviour in complex socio-technical systems. Computer Science Technical Report YCS 390, University of York, 2005.
7. E. Ciapessoni, E. Corsetti, A. Montanari, and P. San Pietro. Embedding time granularity in a logical specification language for synchronous real-time systems. *Science of Computer Programming*, 20:141–171, 1993.
8. E. Corsetti, A. Montanari, and E. Ratto. Dealing with different time granularities in formal specifications of real-time systems. *Journal of Real-Time Systems*, 3(2), 1991.
9. P. Fraisse. *The psychology of time*. New York: Harper and Row, 1963.
10. W. Friedman. *About time: Inventing the fourth dimension*. Cambridge, MA: MIT Press, 1990.
11. E. Hollnagel. *Human Reliability Analysis: Context and Control*. Academic Press, 1993.
12. G. E. Hughes and M. J. Cresswell. *An Introduction to Modal Logic*. University Paperbacks. Routledge, 1968.
13. S.G. Hutchesson and N. Hayes. Technology transfer and certification issues in safety critical real-time systems. In *Digest of the IEE Colloquium on Real-Time Systems*, volume 98/306, April 1998.
14. C.B. Jones, I.J. Hayes, and M.A. Jackson. Deriving specifications for systems that are connected to the physical world. In Cliff B. Jones, Zhiming Liu, and Jim Woodcock, editors, *Formal Methods and Hybrid Real-Time Systems*, volume 4700 of *Lecture Notes in Computer Science*, pages 364–390. Springer Verlag, 2007.
15. M. Joseph, editor. *Real-Time Systems: Specification, Verification and Analysis*. Prentice-Hall, 1996.
16. L. Lamport. Time, clocks, and the ordering of events in a distributed system. *CACM*, 21(7):558–565, 1978.
17. R. Levine. *A geography of time*. New York: Guilford Press, 1997.
18. A. Newell. *Unified theories of cognition*. Harvard University Press, Cambridge, MA, 1990.
19. N.J. Nilsson. Teleo-reactive programs for agent control. *Journal of Artificial Intelligence Research*, 1:139–158, 1994.
20. N.J. Nilsson. Teleo-reactive programs and the triple-tower architecture. *Electronic Transactions on Artificial Intelligence*, 5:99–110, 2001.
21. J.E. Roeckelein. *The concept of time in psychology: A resource book and annotated bibliography*. CT: Greenwood Press, 2000.
22. W. Schneider. Training high-performance skills: Fallacies and guidelines. *Human Factors*, 27(3):285–300, 1985.
23. H.A. Simon. *The Science of the Artificial - 3rd Edition*. MIT Press, 1996.
24. J. M. Spivey. *The Z Notation: A Reference Manual*. Prentice Hall International, second edition, 1992.

A Notation

A partial order “ \preceq ” on some set X is reflexive, transitive, and anti-symmetric.

[X]
$partial_order : \mathbb{P}(X \leftrightarrow X)$
$partial_order = \{ \preceq : X \leftrightarrow X \mid \forall x1, x2, x3 : X \bullet$ $x1 \preceq x1 \wedge$ $(x1 \preceq x2 \wedge x2 \preceq x3 \Rightarrow x1 \preceq x3) \wedge$ $(x1 \preceq x2 \wedge x2 \preceq x1 \Rightarrow x1 = x2) \}$

[X]
$symmetric_rel : \mathbb{P}(X \leftrightarrow X)$
$symmetric_rel = \{ \simeq : X \leftrightarrow X \mid (\forall e, f : X \bullet$ $e \simeq e \wedge (e \simeq f \Leftrightarrow f \simeq e)) \}$