Chapter 4
Time Bands in Systems Structure

Alan Burns and Gordon Baxter
University of York, UK

1 Introduction

One characteristic of complex computer-based systems is that they are required to function at many different timescales (from microseconds or less to hours 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 basic description, and fails to support the separation of concerns that the different timescales of the system facilitate. 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, timescales.

Time is both a means of describing properties of structures and is a structuring mechanism in its own right. To make better use of 'time', with the aim of producing more dependable computer-based systems, it is desirable to explicitly identify a number of distinct *time bands* in which the system is situated. Such a framework enables the temporal properties of existing systems to be described and the requirement for new or modified systems to be specified. The concept of time band comes from the work of Newell[8] in his attempts to describe human cognition. Newell focuses on hierarchical structures within the brain and notes that different timescales are relevant to the different layers of his hierarchy. By contrast, we put the notion of a time band at the centre of any description of system structure. It can then be used within any organisational scheme or architectural form - for they all lead to systems that exhibit a wide variety of dynamic behaviours.

In this work we provide an informal description of a framework built upon the notion of time bands. We consider the properties of bands, how activities within a band can be organised and evaluated, and how relationships between bands can be captured and described. Note that the number of bands required and their actual granularity is system-specific; but the relationships between bands, we contend, exhibit important invariant properties.

2 Motivation and Informal Description

A large computer-based 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 timescales as described above move from the 1ms neuron firing 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 [10] and the development of highly skilful behaviour may take many years. At the organisational and social level, timescales 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.

2.1 Definition of a Band

A band is represented by a granularity (expressed as a unit of time that has meaning within the band) and a precision that is a measure of the accuracy of the time frame defined by the band. System activities are placed in some band B if they engage in significant events at the timescale 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 five minute band, work shifts are changing, meetings are starting, etc. Time therefore has strong discrete properties but is also used to model rates of change etc. For any system there will be a highest and lowest band that gives a system boundary - although there will always be 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.

By definition, all activities within band B have similar dynamics and it may be easy to identify components with input/output interactions or precedence relationships. Within a band, *activities* have duration whilst *events* are instantaneous - "take no time in the band of interest". Many activities will have a repetitive cyclic behaviour with either a fixed periodicity or a varying pace. Other activities will be event-triggered. Activities are performed by agents (human or technical). In some bands all agents will be artificial, at others all human, and at others both will be evident (to an observer or system designer). The relationship between the human agent and the time band will obviously depend on the band and will bring in studies from areas such as the psychology of time[3, 4, 9] and the sociology of time[7].

In the specification of a system, an event may cause a response 'immediately' - meaning that at this band the response is within the granularity of the band. This helps eliminate the problem of over specifying requirements that is known to lead to implementation difficulties [6]. For example, the requirement 'when the fridge door opens the light must come on immediately' apparently gives no scope for an implementation to incorporate the necessary delays of switches, circuitry and the light's own latency.

Events that are instantaneous at band B will map to activities that have duration at some lower band with a finer granularity - we will denote this lower band as C. A key property of a band is the precision it defines for its timescale. This allows two events to be simultaneous ("at the same time") in band B even if they are separated in time in band C. This 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. That is, terms like 'late' are band specific and would not be applied in the latter case. Of course the precision of band B can only be explored in a lower band.

From a focus on band B two adjacent bands are identified. The slower (broader) band (A) can be taken to be unchanging (constant) for most issues of concern to B (or at least any activity in band A will only exhibit a single state change during any activity within band B). At the other extreme, behaviours in (the finer) band C are assumed to be instantaneous. The actual differences in granularity between adjacent bands 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 at 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 well as the system itself manifesting behaviour at many different time bands, the environment will exhibit dynamic behaviour at many different granularities. The bands are therefore linked to the environment at the level determined by these dynamics.

2.2 Behaviour Within a Band

Most of the detailed behaviour of the system will be specified or described within bands. Issues of concurrency, resource usage, scheduling and planning, response time (duration) prediction, temporal validity of data, control and knowledge validity (agreement) may be relevant at any band. Indeed the transfer of techniques from one band to another is one of the motivations for the framework. However, the focus of this paper is on the bands themselves and the relationships between bands, and hence we will not consider in detail these important issues.

We do note however that with human agents (and potentially with artificial learning agents) time itself within a band will play a central role. Time is not just a parameter of a band but a resource to be used/abused within the band. Users will interpret system behaviour from temporal triggers. In particular the duration of an activity will be a source of knowledge and possibly misconceptions, and may be used to give validity (or not) to information, or to infer failure. This use of temporal information to infer knowledge is termed *temporal affordance*[2]. For some bands, agreement (distributed consensus) may depend heavily on such affordances. Plans, schedules or even just routines may give rise to these affordances. They provide robustness, which may be defined into the system but are often developed informally over time by the users of the system (i.e. they are emergent properties). Affordances may be extremely subtle and difficult to identify. Nevertheless the movement of an activity from one band to another (usually a quicker one) may undermine existing affordances and be a source of significant decreased dependability.

Linked to the notion of affordances is that of *context*. A ten minute delay may be a crisis in one context or an opportunity within another. Context will be an issue in all bands but will place a particularly crucial role at the human-centered levels. Context will also play a role in scheduling and planning.

Within a band, a coherent set of activities and events will be observed or planned, usually with insufficient agents and other resources. Robustness and other forms of fault tolerance will also play a crucial role in the description/specification of the behaviour within a band. The specification of some behaviours will require a functional view of time that places ‘time’ at the centre of the design process. To support this process a range of visualisation, modelling and analysis techniques are available including, timed sequence charts, control theory, scheduling analysis, constraint satisfaction, queueing theory, simulation, temporal and real-time logics, timed automata, timed Petri nets, model checking and FMEA (failure modes and effects analysis).

In all bands, a common set of temporal phenomena and patterns of behaviour are likely to be exhibited by the system itself or its environment. For example, periodic (or regular or cyclic) activities, event handling (responding to an event by a deadline), temporal reasoning (planning and scheduling), interleaving and multi-tasking (and other aspects of concurrency), pausing (or delaying), analysis of response (or completion) time, deadline driven activities, and various aspect of dynamic behaviour such as rates of change. Whilst evident in all bands, these phenomena are not identified using the same terminology in the various time bands of interest (i.e in the technical, psychological and sociological literature). The development of an agreed collection of guide words within the framework would therefore help link temporal issues with other significant phenomena within a specific band (e.g. terms such as temporal memory, event perception etc. within a 'psychological' band).

We also note that the vocabulary usually associated with temporal issues (e.g. late, too soon, on time, simultaneous, instantaneous, immediate, before, never, having enough time, running out of time, plenty of time, etc) can be given quite specific meanings if they are made band specific. For example, in a human-centred band an electronic spreadsheet responds immediately. Of course at a much lower level band considerable activities are needed to furnish this behaviour. Making the vocabulary of requirements in systems explicitly band-specific will remove some of the misconceptions found with regard to timing issues in such documents.

Finally, we emphasize that the framework is not reductionist. Lower bands contain more detail about individual events. Higher bands contain information about the relationships between activities in a more accessible form. Emergent properties will be observed within a band. The motivation for the framework is to be able to describe these properties, and where necessary link them to more primitive actions at a lower band.

2.3 Behaviour Between Bands

To check the coherence of a description, or the consistence of a specification, for a complex socio-technical system, requires behaviours between bands to be examined. This involves two issues:

1. the relationship between the bands themselves, and
2. the mapping of activities and events between bands.

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 5 minutes is easily linked to a band with a granularity of 10 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 5 minutes 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 2.55 and 3.05.

Granularity can however give rise to a more complex link. For example a band with a granularity of a month can be linked to a band with granularity of a day by linking a month to between 28 and 31 days. Here precision is exact; both bands have the same notion of accuracy about the time reference.

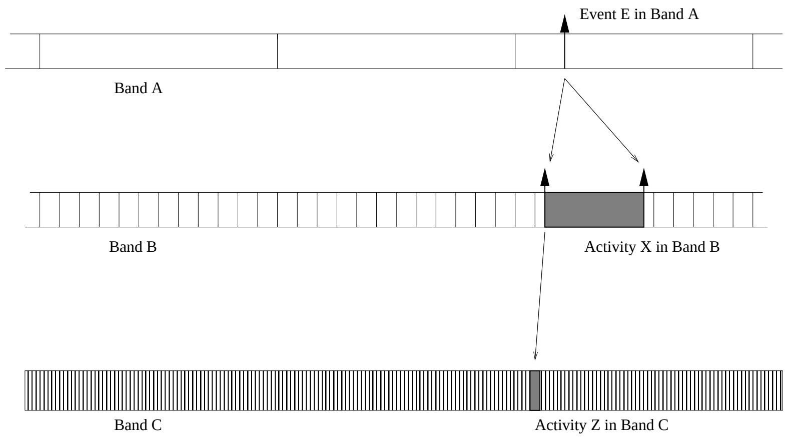

Figure 1: Time Band Example

The mapping of actions between bands is restricted to: event to event, or, event to activity relations. So an event in some band can be identified as being coupled to (implemented by) an event or activity in some other band. A specific named activity exists in one, and only one, band. But for all activities there are events that as defined to note the start and end of an activity - these events can be mapped to finer bands. Moreover the whole activity can be seen as an event in a broader band. Figure 1 illustrates three bands (A, B and C) with an event E in band A being mapped to activity X in band B. The start and end events of this activity are then associated with activities in band C.

To exercise these abstracts, consider the planning of a university curriculum. When planning courses on a term by term basis, a lecture is an event. When planning room allocations, a lecture becomes an activity of duration one or two hours (with precision 5 minutes). When planning a lecture, each slide is an event (with an implicit order). When giving a lecture each slide is an activity with duration. This description could be given in terms of a number of bands and mappings of events to activities in finer bands. Note when focusing on the band in which slides have duration it is not possible or appropriate to consider the activities in higher bands that represent whole courses or semesters. The time bands therefore correctly helps to separate concerns. Students may learn that the time spent on a slide implies importance (at least in terms of likelihood of turning up in an exam). This is an example of a temporal affordance. Also illustrated by this situation is the difference between planned behaviour (as one moves down the time bands) and emergent properties that enable students to structure the knowledge and understanding they have obtained in many different ways during their progression through their degree course.

To return to the crucial issue of coherence and consistence between bands, the proposed framework facilitates this by making explicit the vertical temporal relationships between bands. Specifically, it becomes possible to check that the temporal mapping between event E in band A with activity X in band B is consistent with the bounds on the relationship identified between bands A and B. Moreover this consistency check can be extended to ordered events and causality (see next section). So, to give a simple example; a lecture of 10 slides each with duration 5 minutes (precision 1 minute) cannot implement the lecture event (as this was mapped to an activity with duration one hour and precision 5 minutes).

2.4 Precedence Relations, Temporal Order and Causality

At the time bands of 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. But there are contexts in which *order* is a more natural way of describing behaviour [1, 5](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. 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 clock in each distinct time zone.

There is of course a strong link between temporal order (i.e. time stamped events and activities) and precedence relations. If P is before Q then in no band can a time stamp for P be after that of Q. 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 P and Q is sufficiently large in the band of interest.

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.

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”. As noted earlier within the band of human experience this can be taken to be ‘immediate’ and modelled as an event. At a lower band a number of electromechanical activities will be needed to be described that will sense when the door is open and enable power to flow to the light.

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, as noted above, order and hence causality is not necessarily maintained as one moves down through the bands. This is a key property of the framework, and implies that where order is important then proof must be obtained by examining the inter-band relationship (as discussed above).

2.5 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 behavioural description of a system. System activities are always relative to (defined within) a band, and have duration of one or more ticks of the band’s granularity. Events in band B take no time in that band, but will have a correspondence with activities within a lower band. It follows that a number of events can take place “at the same time” within the context of a specified band. Similarly responses can be “immediate” within a band.

Precedence relations between activities and events are an important part of the framework and allow causal relations to be defined without recourse to explicit references to time. Moreover they can be used to define clock tick events within a band, and hence link other events to the absolute time of the band.

We require all time bands to be related to one another but do not require a strict perfect mapping. Each band, other than the lowest, will have a precision that defines (in a lower band) the tolerance of the band. However within these constraints we do need to be able to show that system descriptions at different bands are consistent. For this to be possible a formal description is required.

3 Case Study

Omitted in this abridged version.

4 Conclusion

In this chapter we have argued that complex systems exhibit behaviour at many different time levels and that a useful aid in structuring, describing and specifying such behaviour is to use time bands. Viewing a system as a collection of 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. Indeed the imprecision may be large in social systems and be a source of dependability (robustness).

Acknowledgements

The time band formulation owes much to the input of Colin Fidge and Ian Hayes. This particular description has benefited from comments of Denis Besnard and Cliff Jones.

References

[1] J. Allen. Towards a general theory of actions and time. *Artificial Intelligence*, 23:123–154, 1984.

[2] A.P. Conn. Time affordances: the time factor in diagnostic usability heuristics. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 186–193, 1995.

[3] P. Fraisse. *The psychology of time*. New York: Harper and Row, 1963.

[4] W. Friedman. *About time: Inventing the fourth dimension*. Cambridge, MA: MIT Press, 1990.

[5] E. Hollnagel. *Human Reliability Analysis: Context and Control*. Academic Press, 1993.

[6] 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.

[7] R. Levine. *A geography of time*. New York: Guilford Press, 1997.

[8] A. Newell. *Unified theories of cognition*. Harvard University Press, Cambridge, MA, 1990.

[9] J.E. Roeckelein. *The concept of time in psychology: A resource book and annotated bibliography*. CT: Greenwood Press, 2000.

[10] W. Schneider. Training high-performance skills: Fallacies and guidelines. *Human Factors*, 27(3):285–300, 1985.