

EVALUATING TIMEBANDS AS A TOOL FOR STRUCTURING THE DESIGN OF SOCIO-TECHNICAL SYSTEMS

Gordon Baxter¹, Alan Burns² and Kenneth Tan³

¹ *Department of Psychology, University of York, Heslington, York YO10 5DD*

² *Department of Computer Science, University of York, Heslington, York YO10 5DD*

³ *Department of Pediatrics, McMaster University, Hamilton, Ontario, Canada*

Structuring the design of a socio-technical system so that it fully meets its requirements is a complex problem. Typically, the system is described using hierarchical structures, but time is only represented as a single flat physical phenomenon. We have developed a framework that allows systems to be partitioned into a hierarchical series of time bands, such that the temporal properties can be adequately represented and exploited. We have used a simple scheduling model to evaluate the framework using sample scenarios from a case study in a neonatal unit. The results highlight: the importance of finding the right level of abstraction when identifying system activities; the need to ensure that appropriate precedences are allocated to activities that are part of larger procedures; and offer some support for the separation of concerns between bands.

Introduction

The performance of socio-technical systems is the result of a combination of human, technological and organisational factors. Structuring the design of these systems so that they fully meet their requirements is a complex problem. It is important to solve this problem, though, for domains such as transportation and medicine, where the timing of actions is critical to system dependability, and the systems have to simultaneously function at many different time scales (from microseconds, or less, to hours, or more). Whilst these systems are typically described using hierarchical structures, time is often represented as a single flat physical phenomenon. This abstraction forces different temporal notions onto the same flat description, and fails to support the separation of concerns that the system's different time scales facilitate. An alternative view, adopted here, is that systems can be partitioned into a hierarchical series of time bands (Burns & Baxter, 2006). (For a more formal definition of time bands see Burns et al., 2005)

The notion of time bands is summarised below. We then briefly consider the idea of modelling actions within Human-Computer Interaction (HCI). This concept of modelling is then used in applying the time bands framework to some simple scenarios based on

data collected during a case study of a neonatal intensive care unit. Finally we discuss the results of the modelling and indicate some areas of future work.

Time bands

We have developed a framework that explicitly identifies a number of distinct *time bands* in which the system under consideration is situated. This notion of time bands is partly inspired by Newell (1990) who noted that the actions associated with HCI fall into two time bands: the cognitive band (100 ms to 10 s) and the rationality band (10^2 to 10^4 s).

The *time band* abstraction is an attempt to derive an engineering framework based on natural temporal phenomena. Most formulations that attempt to identify time granularity do so by mapping all activities onto the finest granularity in the system. Here we assume that a system is composed of a partially ordered finite set of time bands.

A band is primarily represented by a *granularity*. System activities are placed in some band **B** if they engage in significant work at the time scale represented by **B**. For any system there will be a highest and lowest band that gives a temporal system boundary. By convention the lower bands have the finer granularity.

The second defining characteristic is *precision*, which is the degree of tolerance of a time statement within the band. A band's precision can only be measured in a lower band when the units of the precision can be articulated within the granularity of the lower band.

Within each band, there are *activities*, which have duration, and *events*, which are instantaneous, in that they take no time in the band of interest. Activities are performed by *agents* (human or artificial). There can be a mix of different types of agents within the same band, and some agents will perform activities in more than one band.

In the specification of a system, an event may cause a response *immediately*, that is, within the granularity of the band. This helps eliminate the problem of overspecifying requirements that is known to lead to implementation difficulties. Making the term *immediate* band specific enables a finer granularity band to include the necessary delays, latencies and processing time required to support the immediate behaviour at the higher band. The obligation on this lower band is to deliver the behaviour within the precision of the higher band. Hence, events that are instantaneous in some band may map to activities that have duration at some lower band. The key relationship that enables the framework to support system decomposition and modelling, is that within any band, activities in lower (faster) bands are assumed to be instantaneous and the state of activities in higher (slower) bands are assumed to be unchanging.

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 prediction, temporal validity of data, control and knowledge validity may be relevant in any band. Time is not just a parameter of a band but a resource to be used (or even 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 validate information, or to infer failure.

Modelling actions

There is a tradition of using engineering models within HCI. In most cases, however, these models focus on low-level interactions, such as pressing keys and pushing buttons.

We know, however, that actions that occur at higher levels can affect lower levels. So, for example, local operating procedures may require operators to follow particular sequences of actions in carrying out some predefined task.

We have adopted a simple scheduling approach to conduct a preliminary evaluation of the time band framework, by modelling scenarios from a case study of a hospital neonatal unit (Baxter et al., 2005). Rather than focusing on single person, multi-task situations, however, we have been looking at the utility of time bands for describing multi-person, multi-task situations.

Applying the time bands framework to neonatal intensive care

In the neonatal intensive care unit (NICU), staff routinely use technology in treating premature babies (by delivering drugs and food, assisting breathing, monitoring the baby's vital signs and so on). One of the problems often faced by these babies is respiratory distress syndrome. This is treated by a combination of medication and mechanical ventilation. The aim is to ensure that the baby's blood gases remain within predefined limits whilst it recovers from the self-regulating disease. The blood gases are continuously monitored, and if they go out of range, adjustments are made to the settings on the mechanical ventilator that helps the baby to breathe. It is crucial that these adjustments are made in a safe and timely manner.

Time bands in the neonatal unit

A preliminary analysis of the neonatal unit suggested four identifiable time bands:

- Future Planning (Granularity: one week; precision: one day).
- Ward Organisation (Granularity: half an hour; precision: 10 minutes).
- Clinical Procedures (Granularity: five minutes; precision one minute).
- Baby Dynamics (Granularity: one second; precision ten milliseconds).

Note that we originally collected the data for the purposes of identifying timing issues and how they affected the dependability of the delivery of neonatal care. In other words, we did not collect data specifically so that it could be modelled using time bands.

The model

We used a simple scheduling model, adapted from that of Bratko (2001). The model was written in Scistus Prolog using Constraint Logic Programming over Finite Domains (Carlsson et al., 1997). The model takes a set of activities, a set of agents, a set of pairwise activity preferences, and a set of activity allocations (indicating which agents can perform which activity) and, by applying a set of constraints, determines the optimum time and sequence of the activities.

Scenario 1: Modelling the clinical procedures band

We began by developing a simple, but realistic scenario. We initially only modelled a single time band, because we wanted to test the utility of the scheduling model.

In our scenario there were 4 nurses and 3 junior doctors caring for 8 babies. In addition there were 2 registrars and 1 consultant available locally, all of whom could be called in to assist. Alarms sound for four of the babies, indicating a problem that requires investigation. For one baby, changes to the ventilator settings are required to bring the blood gases back within limits. These changes do not fix the problem, however, so a registrar is called in and makes further changes. The problem remains unresolved, so the

consultant is called in, and makes changes to the ventilator settings that finally fix the problem. In parallel with this acute situation, one of the other babies has a routine clinical procedure carried out.

Initially the model indicated that the total time to perform all the activities was 1500 s (25 minutes), which was much shorter than anticipated, given that there were three periods of 20 minutes that were needed to allow the ventilator changes to take full effect. Closer inspection of the data showed that the model had allocated several activities in parallel to different members of staff, when they needed to be executed serially.

The underlying problem was that the level of abstraction for the *medical intervention* activity was too coarse. These interventions differ qualitatively, and are based on the level of expertise and experience of the member of staff. This explains why junior doctors often only deal with simple cases, and pass on more complex cases to the registrars or consultants, as appropriate. After revising the model to reflect these differences, and introducing precedence relations to ensure that activities were completed in the correct order, the time increased to 5160 s (86 minutes), in line with our initial expectations.

Scenario 2: Modelling multiple bands

The second scenario revolved around the need for an X-ray to be taken, a procedure that takes place in the Ward Organisation Band. One of the doctors calls in the radiographer to take an X-ray. The radiographer comes along to the NICU, takes the X-ray picture, and goes away to develop it before returning to deliver it to the doctors in the NICU.

After running this scenario on its own to ensure that the time taken and the sequence of activities was correct, we modelled the two scenarios in parallel. Initially there was an apparent anomaly in the allocation of activities to bands. Monitoring the data after an intervention takes 20 minutes, and is located in the Clinical Procedures Band; developing an X-ray picture also takes 20 minutes but resides in the Ward Organisation Band. We resolved the issue by making the former a single activity in the Ward Organisation Band, and adding appropriate events to that band which map onto the activities in the Clinical Procedures Band. In this way, it becomes possible to avoid the potential conflicts of agents being allocated to perform activities simultaneously in adjacent time bands. The outcome is that bands have to be modelled separately, otherwise agents can be allocated to both an event in one band and its corresponding activity in the next lower band.

The time to perform all the activities in the Ward Organisation Band was 5160 s (40 minutes), which was in line with our expectations. The time to perform all the activities in the Clinical Procedures Band had now reduced to 1560 s. The total time for both scenarios is 5160 s (81 minutes). This is the same as the time from scenario 1, where only one band was modelled. The reason for this is that in the second (X-ray) scenario, the procedure can be carried out in parallel with all of the other activities by using free agents (the third junior doctor and the radiographer).

Summary and future work

Although this work is fairly preliminary, it has yielded some useful insights. These will be used to refine our thinking about the time bands framework, and to help shape how we go about collecting and analysing data that we can model using time bands in the future.

The level of abstraction of the activities is very important. In our first analysis of the data from the NICU, we simply grouped all the medical interventions into one type of activity. It was only when we started to model the data that we fully appreciated the need

to distinguish between routine (or planned) interventions, and acute interventions, and that the latter can be subdivided into qualitatively different activities that are reliant on the experience and expertise of a particular type of agent. This more fine-grained categorisation is particularly important when it comes to determining the temporal ordering of activities.

Whilst time bands are populated with activities (and events), it is clear that the activities are often (but not always) steps in a much larger procedure. These activities may not be carried out by the same individual within a particular procedure. In the Clinical Procedures Band modelling, for example, we found that dealing with an acute incident can involve nurses, junior doctors, registrars and consultants. Furthermore, there is a degree of non-determinism about some of these procedures. So, when a junior doctor changes the ventilator settings, it is only when the full response to the changes become apparent (20 minutes later) that the success of the procedure can be evaluated (and the procedure terminated), or whether it needs to be continued by calling in a registrar.

The fact that when multiple bands are involved, the bands have to be modelled individually supports the idea of separating the concerns for the different bands. There may be links between events in one band and activities in a lower band, and it is important to make sure that these do not both allocated in the activity schedule for the system. In addition, there is no apparent need for the higher band to be concerned with events in the lower band, because these events will usually have a duration that is much shorter than the precision of the higher band.

There is still much work that needs to be done to fully flesh out the time bands framework. Perhaps the most obvious next step, is to develop a model using three time bands, to see if adding more time bands yields any further insights (or complications!). On the basis of our work so far, though, it appears as though the time bands framework does offer a valid way for organising the structure of socio-technical systems.

References

- Baxter, G. D., Filipe, J.-K., Miguel, A. and Tan, K. 2005, The effects of timing and collaboration on dependability in the neonatal intensive care unit. In F. Redmill and T. Anderson (eds.) *Constituents of modern system -safety thinking: Proceedings of the thirteenth safety-critical systems symposium*, (Springer Verlag, Berlin, Germany), 195-210
- Bratko, I., 2001, *Prolog programming for artificial intelligence*, Third edition, (Addison-Wesley, London, UK)
- Burns, A. and Baxter, G. 2006, Time bands in system structure. In D. Besnard, C. Gacek and C. Jones (eds.) *Structure for Dependability: Computer-based Systems from an Interdisciplinary Perspective*, (Springer, London, UK), 74-90
- Burns, A., Hayes, I. J., Baxter, G. and Fidge, C. J. 2005, *Modelling Temporal Behaviour in Complex Socio-Technical Systems*. (University of York - Department of Computer Science: York)
- Carlsson, M., Ottosson, G. and Carlson, B. 1997, An open-ended finite domain constraint solver. In H. Glaser, P. Hartel and H. Kuchen (eds.) *Proc. Programming Languages: Implementation, Logics and Programs.*, (Springer-Verlag, Berlin), 191-206
- Newell, A., 1990, *Unified theories of cognition*, (Harvard University Press, Cambridge, MA)