Timed Communicating Object Z

Brendan Mahony

Information Technology Division
Defence Science and Technology Organisation (DSTO)
Brendan.Mahony@dsto.defence.gov.au

Jin Song Dong

Mathematical and Information Sciences
Commonwealth Scientific and Industrial Research Organisation (CSIRO)
jinsong.dong@cmis.csiro.au

February 1998

Abstract

This paper describes a timed, multi-threaded object modeling notation for specifying real-time, concurrent, and reactive systems. The notation Timed Communicating Object Z (TCOZ) builds on Object-Z's strengths in modeling complex data and algorithms, and on Timed CSP's strengths in modeling process control and real-time interactions. TCOZ is novel in that it includes timing primitives, properly separates process control and data/algorithm issues and supports the modeling of true multi-threaded concurrency. TCOZ is particularly well suited for specifying complex systems whose components have their own thread of control. The expressiveness of the notation is demonstrated by a case study in specifying a multi-lift system that operates in real-time.

1 Introduction

Many formal specification and design notations have tended to concentrate either on data modeling and algorithmic concerns (eg. Z, VDM, etc.) or else on process control concerns (eg. CSP, CCS, StateCharts, etc.). Complex systems often have intricate system states and process control structures involving concurrency and real-time interactions. To formalise such systems, it is necessary to have a notation which is able to capture both the data/algorithmic issues and the process behaviour issues in a smoothly integrated, but also highly structured and modular, manner. In consequence, the blending of state-modeling and process languages has become an active area of research [19, 43, 44, 21].

Object-Z [17] is an object oriented extension of the Z formal specification language. Z is a model-oriented specification language with powerful features for describing complex data structures and their operations. Object-Z improves the clarity of large Z specifications through enhanced structuring. However, the process semantics of Object-Z mean that: process execution is single threaded; operations are atomic, there is no notion of the duration of operations; and process control logic is tightly coupled with class structure. Therefore, it is difficult to use Object-Z to model concurrent real-time reactive systems.

Timed CSP [40] is an extension of Hoare's Communicating Sequential Processes (CSP) notation. It builds on CSP's strengths in modeling process control issues, such as concurrency and synchronisation, by adding primitives for modeling real-time issues. However, CSP has only the most rudimentary mechanisms for modeling data and algorithmic issues and it is cumbersome to capture the state of a complex system.

This paper describes an integration of Object-Z and Timed CSP, called Timed Communicating Object Z (TCOZ), and presents a case study on using TCOZ to specify a real-time multi-lift system. TCOZ builds on the respective strengths of the Object-Z and Timed CSP notations in order to provide a single notation for modeling both the state and process aspects of complex systems. The notion of blending Object-Z with CSP has been suggested independently by Fischer [19] and Smith [43]. The most obvious novelty of TCOZ is that it is built on Timed CSP and includes primitives for treating timing issues, however in addressing the issue of time it has been necessary to make several inovations which impact positively even on the treatment of 'untimed systems'. TCOZ adopts a finer grain of atomicity than either Fischer or Smith. Operations are considered to represent a sequence of (unspecified) update events, rather than to constitute atomic events in themselves. This opens the possibility of treating operation composition and refinement in TCOZ, including the introduction of multithreaded concurrency at the operation level. TCOZ adopts an explicit mechanism for enabling operations (and indeed arbitrary processes) which is distinct from the operation definition itself. This increases the potential for reuse of operation specifications and allows the notions of operation and process refinement to be reconciled. TCOZ adopts the CSP channel-based communications paradigm in its full generality and enhances it by the introduction of a novel network topology operator that allows the communications interfaces of complex TCOZ processes to be visualised through simple network-topology graphs. This improves decoupling of class definitions by simplifying the interfaces between objects. For the most part, these topics can be touched on only briefly in this paper and they will be the subjects of future more detailed correspondences.

The TCOZ notation has been briefly described and exercised in introductory papers by these authors [34, 35]. This paper combines elements of these papers, but describes the notation and its use in greater detail. Important issues such as formal syntax, class inheritance, object composition, and active/passive objects are explained for the first time.

The remainder of the paper is organised as follows. In Section 2, Object-Z and Timed CSP notations are briefly introduced. The advantages and disadvantages of the two notations in modeling timing, concurrency, complex data and algorithmic aspects are demonstrated using a common example, a timed collection. In Section 3, the blended notation, TCOZ, is introduced and the timed collection example used to show how it uses the strengths of the individual notations to address their respective weaknesses. In Section 4, the case study on specifying a real-time multi-lift system is presented. In Section 5, a discussion of related work is presented. Finally, a syntax for TCOZ is presented as Appendix A.

2 Object-Z and Timed CSP

The common example of a generic timed collection is used through out this section and the next section to to illustrate the differences between and to demonstrate the advantages and disadvantages of the Object-Z, Timed CSP, and TCOZ notations respectively.

2.1 Generic timed-collection example

The generic timed-collection denotes a collection of elements of type X with a time stamp. Operations are allowed to add elements to and delete elements from the collection. When deleting an element from the collection, the oldest element should be removed and output to the environment. The collection has the following timing properties. Firstly, that it takes a small but non-zero time (t_a and t_d respectively[1]) to update the internal state during a *add* or *delete* operation. Secondly, each element of the collection becomes *stale* if it is not passed on within t_o time units of being added to the collection. Stale elements should never be passed on, but are instead purged from the collection upon becoming stale. The *purge* operation has a duration of t_p.

[1] For ease of presentation (especially in the Object-Z and Timed CSP versions) we adopt exact timing constraints in the timed-collection example.

2.2 A model of time

In this paper, all timing information is represented as real valued measurements in *seconds*, the SI standard unit of time [27]. Describing time and other physical quantities in terms of standard units of measurement is an important aspect of ensuring the completeness and soundness of specifications of real-time, reactive, and hybrid systems. In order to support the use of standard units of measurement, extensions to the Z typing system suggested by Hayes and Mahony [24] are adopted. Under this convention, time quantities are represented by the type

$$\mathbb{T} == \mathbb{R}\,\mathsf{s},$$

which represents real-valued time measured in seconds. Time literals consist of a real number literal annotated with a symbol representing a unit of time. For example, $3\,\mu\mathsf{s}$ is a literal representing a period of three microseconds. All the arithmetic operators are extended in the obvious way to allow calculations involving units of measurement.

The timing constants associated with the timed-collection example are introduced via axiomatic definitions.

$$t_a, t_d, t_p, t_o : \mathbb{T}$$

2.3 Object-Z

The main Object-Z construct is the *class* definition. A class is a template for *objects* of that class: for each object of a class, the object's states are instances of the class' state schema and the object's state transitions are instances of the class' operation schemas. An object is said to be an instance of a class and to evolve according to the definitions of its class.

Since Object-Z has no standard conventions for handling timing and process control issues, it is necessary to model these issues explicitly in the class state. One such approach is Timed Object-Z [10], which incorporates ideas from various Z-based approaches for specifying real-time requirements [18, 36]. The Timed Object-Z approach consists of two conventions. Firstly, environmental factors are modeled as functions of time and are included in the system state. In the timed-collection example, the environment is modeled by functions *left* and *right*. These functions represent the participation of the environment in *Add* and *Delete* operations respectively. The second extension is to include a global real-time clock, conventionally represented by a distinguished state attribute *now*. The clock may only be updated during an operation and the next operation must start as soon as the previous one is finished. In order to model the time taken between operations it is thus necessary to introduce a *Wait* operation.

2.3.1 Timed-collection in Object-Z

The use of Timed Object-Z is illustrated by the *TimedCollection* class presented below.

TimedCollection$[X]$

$$left, right : \mathbb{T} \pfun X$$
$$mems : \mathbb{F}(\mathbb{T} \times X)$$
$$\Delta$$
$$now : \mathbb{T} \qquad \text{[current time]}$$
$$t_i, idle : \mathbb{T} \qquad \text{[idling parameters]}$$
$$(t, oldest) : \mathbb{T} \times X \qquad \text{[oldest item]}$$

$$mems \neq \emptyset \Rightarrow (t, oldest) \in mems \land t = min\,\mathrm{dom}\,mems$$
$$\textbf{let}\ rm = \{\tau : \mathrm{dom}(left \cup right) \mid \tau \geq now \bullet \tau - now\} \cup \mathrm{dom}\,mems \bullet$$
$$rm \neq \emptyset \Rightarrow idle = min(rm)$$

INIT

$$mems = \varnothing \land now = 0\,\mu\mathsf{s} \land t_i = 0\,\mu\mathsf{s}$$

Add

$\Delta(mems)$
$e? : X$

$now \in \operatorname{dom} left \land left(now) = e?$
$now' - now = t_a \land t_i' = 0\,\mu\mathsf{s}$
$mems' = ps(t_i + t_a, mems) \cup \{(t_o, e?)\}$

Delete

$\Delta(mems)$
$e! : X$

$mems \neq \varnothing$
$now \in \operatorname{dom} right \land right(now) = e!$
$now' - now = t_d \land t_i' = 0\,\mu\mathsf{s}$
$e! = oldest$
$mems' = ps(t_i + t_d, (mems \setminus \{(t, oldest)\})$

Purge

$\Delta(mems)$

$t_i \in \operatorname{dom} mems$
$now' - now = t_p \land t_i' = 0\,\mu\mathsf{s}$
$mems' = ps(t_i + t_p, mems)$

Wait

$idle > 0\,\mu\mathsf{s}$
$now' - now = idle$
$t_i' = t_i + idle$

The generic function *ps* (purge stale) is defined as

$[X]$

$$ps : (\mathbb{T} \times \mathbb{F}(\mathbb{T} \times X)) \to \mathbb{F}(\mathbb{T} \times X)$$

$$ps(t, \varnothing) = \varnothing$$
$$\forall\, t : \mathbb{T};\ s : \mathbb{F}(\mathbb{T} \times X) \bullet ps(t, s) = \{(t_o, e) : s \mid t_o > t \bullet (t_o - t, e)\}$$

e.g. $ps(2\,\mu\mathsf{s}, \{(1\,\mu\mathsf{s}, a), (3\,\mu\mathsf{s}, b), (7\,\mu\mathsf{s}, c)\}) = \{(1\,\mu\mathsf{s}, b), (5\,\mu\mathsf{s}, c)\}$.

The first schema of a class is called the state schema. It describes the various state *attributes* and the class *invariant*. The class attributes are divided into *primary* attributes and *secondary* attributes, which appear below a Δ separator placed in the declaration section of the state schema. The important aspect of secondary attributes in the context of the *TimedCollection* classes is the fact that they are subject to change by every operation. For a detailed discussion of secondary attributes see Dong *et al* [12]. The predicate below the line in the state schema is called the class invariant. It describes the state properties that must be established initially and preserved by every operation.

In the *TimedCollection* class, the primary attribute *mems* denotes a finite set of time-stamped elements of the generic type X. The other primary attributes are the *left* and *right* environment variables. Since these variables represent interactions with the environment, they are not subject to change by any of the class operations.

In the *TimedCollection* class, the secondary attributes are *now*, denoting the current time; and t_i, denoting the idle time since the completion of the last operation; the tuple $(t, oldest)$ denoting the oldest element being in the collection, when it is non-empty; and the time variable *idle*, denoting the time until the next operation.

The INIT schema describes the allowed initial values for the class attributes, in this case *mems*. The initial schema implicitly includes the state schema, so that the initial state includes all the class attributes and satisfies the class invariant. The timed-collection initially contains no elements.

The remaining four schemas are operation schemas which describe the allowed state transitions for the class. The declaration parts of operation schemas may include a Δ-list of those (primary) attributes whose values may change. By convention, any primary attribute not appearing in the Δ-list may not change value. The values of the secondary attributes are always subject to change. Every operation schema implicitly includes the state schema in un-primed form (the state before the operation) and primed form (the state after the operation).

The timing behaviour of the *TimedCollection* requires careful explanation. The *Add* operation may occur only when there is an item available on *left*. It updates timeouts in the existing collection and purges any stale items (this is described in the *ps* funtion definition), and adds the new item with the maximum timeout of t_o. This

activity must take exactly t_a and the idle time is set to $0\,\mu\mathsf{s}$. The *Delete* operation is enabled only when the *right* environment is willing to accept the oldest item, it communicates the item and then deletes it from the collection. The *Purge* operation is envoked when no communication is possible before the first item goes stale ($t_i \in \mathrm{dom}\, mems$). Each of these operations also implicitly updates the *idle* attribute according to the requirement in the data invariant. The *idle* attribute always records the amount of idling required before the next action is enabled and the *Wait* operation simply consumes this idle time.

2.3.2 Transition system interpretation

Object-Z has a three-stage semantics. The various operations in a class are given a standard Z semantics, then this is used to develop a transition-system semantics, which then determines an event-based process semantics [42]. The Z operation semantics is best viewed as describing a relation between initial and final states for each operation. The operations of a given class thus form a named collection of relations, which determines a transition system in which a given operation may fire exactly when its Z precondition is satisfied. The Z precondition of an operation schema describes the initial states for which there exists some final state satisfying the schema predicate. The process model for the class consists of all the sequences of operations/events which can be performed by objects of the class.[2]

For example, the *TimedCollection* object starts with *mems* empty then evolves by successively performing either *Add*, *Delete*, *Purge*, or *Wait* operations. This is sometimes expressed semi-formally by an equation such as

$$TCbeh \mathrel{\widehat=} (Add \Box Delete \Box Purge \Box Wait);\ TCbeh.$$

Here *TCbeh* represents the behaviour of the *TimedCollection*, $(_ \Box _)$ represents a choice between operations, and $(_;\ _)$ represents sequential composition.[3] The choice of which operations are enabled at each point is determined by which preconditions are satisfied by the current state. As an example of a precondition calculation consider the *Purge* operation. The Z precondition is defined to be

$$\exists\, mems' : \mathbb{F}\,(\mathbb{T} \times X);\ now', t_i', idle' : \mathbb{T};\ (t, oldest) : \mathbb{T} \times X \bullet Purge.$$

By expanding the predicate part of the *Purge* schema and simplifying, it can be shown that this is equivalent to

$$\mathrm{pre}\ Purge \quad now - t_i \in \mathrm{dom}\, mems$$

Thus the *Purge* operation may only occur when the oldest member of the timed-collection has expired.

This entwining of behavioural control matters with algorithmic matters creates unnecessary complexity in the design process and fails to promote a clear separation of concerns. For example, in order to ensure that operations occur in some desired order the designer must painstakingly craft the preconditions of all the operations in a class so as to ensure the desired interactions and may even need to add unnecessary process state in order to represent control state. Since there is no way to progress time except through the action of an operation, it is necessary to introduce pseudo-operations such as *Wait* whose sole purpose is to make sure that there is something to do at each point in time and deep reasoning is required to demonstrate that time does indeed always progress. Perhaps most inconvenient, is the fact that this use of preconditions to control the sequencing of transitions is incompatible with Z's standard algorithmic refinement techniques. Recall that operation-refinement in Z consists of either weakening the precondition of an operation or else strengthening its postcondition. Under this convention, weakening an operation's precondition would result in it being enabled more often, thus playing havoc with the process control structure of the original specification. For example, if the precondition for the *Purge* operation was weakened it would be possible for it to occur either before or after the expiry timestamp of data element, a result completely at odds with the purpose of the timestamp. In Object-Z this conflict is resolved by abandoning Z algorithmic refinement, a decision that effectively divorces Object-Z from both the theory and practise of program design and development using Z.

[2] Smith's semantics also includes *ready* sets which record the enabled events at each step.

[3] Note the clear influence of CSP.

2.3.3 Summary

The *TimedCollection* class describes the data state (*mems*) and operations of the timed collection well. However, Timed Object-Z requires interactions with the environment and the progress of time to be micro-managed in an intrusive manner. The proliferation of additional attributes required to deal with process control and time result in significant over specification of the system. All too frequently deep reasoning is required to comprehend the subtle and complex interplays between operations and environment.

2.4 Timed CSP

Timed CSP [40] extends the well-known CSP (Communicating Sequential Processes) notation of Hoare [26] with timing primitives. CSP is an *event* based notation primarily aimed at describing the sequencing of behaviour within a process and the synchronisation of behaviour (or *communication*) between processes. Timed CSP extends CSP by introducing a capability to quantify temporal aspects of sequencing and synchronisation.

CSP adopts a symmetric view of process and environment. Events represent a co-operative synchronisation between process and environment. Both process and environment may control the behaviour of the other by *enabling* or *refusing* certain events or sequences of events. Although CSP semantics are symmetric with respect to process and environment, we find it helpful in the following to use the words *request* and *block* as synonyms for enable and refuse respectively when referring to the environment.

2.4.1 Process primitives

A process which may participate in event a then act according to process description P is written

$$a@t \rightarrow P(t).$$

The event a is initially enabled by the process and occurs as soon as it is requested by its environment, all other events are refused initially. The event a is sometimes referred to as the *guard* of the process. The (optional) timing parameter t records the time, relative to the start of the process,[4] at which the event a occurs and allows the subsequent behaviour P to depend on its value.

The second form of sequencing is process sequencing. A distinguished event $\checkmark$ is used to represent and detect process termination. The sequential composition of P and Q, written $P; Q$, acts as P until P terminates by communicating $\checkmark$ and then proceeds to act as Q. The termination signal is hidden from the process environment and therefore occurs as soon as enabled by P. The process which may only terminate is written SKIP.

The parallel composition of processes P and Q, synchronised on event set X, is written

$$P \,|[\,X\,]|\, Q.$$

No event from X may occur in $P \,|[\,X\,]|\, Q$ unless enabled jointly by both P and Q. When events from X do occur, they occur in both P and Q simultaneously and are referred to as *synchronisations*. Events not from X may occur in either P or Q separately but not jointly. For example, in the process described by

$$(a \rightarrow P) \,|[\,a\,]|\, (c \rightarrow a \rightarrow Q)$$

all a events must be synchronisations between the two processes. Since a is not enabled initially by the right hand process, a cannot occur in the left hand process until the right hand process has performed a c event and the a event becomes enabled in both processes.

In an asynchronous parallel combination

$$P \,|||\, Q$$

[4] This may be non-zero because the process must wait until the event is requested by its environment.

both components P and Q execute concurrently without any synchronisations.

Diversity of behaviour is introduced through two choice operators. The external choice operator allows a process a choice of behaviour according to what events are requested by its environment. The process

$$(a \rightarrow P) \,\Box\, (b \rightarrow Q)$$

begins with both a and b enabled. The environment chooses which event actually occurs by requested one or the other first. Subsequent behaviour is determined by the event which actually occurred, P after a and Q after b respectively. When the range of choices is large (possibly infinite), external choice may be written in an intentional form,

$$\Box\, a : A \bullet P(a),$$

which allows the environment to choose any event a from a set A and subsequent behaviour is determined by $P(a)$.

Internal choice represents variation in behaviour determined by the internal state of the process. The process

$$a \rightarrow P \sqcap b \rightarrow Q$$

may initially enable either a, or b, or both, as it wishes, but must act subsequently according to which event actually occurred. The environment cannot affect internal choice. Again an intentional form is allowed.

An important derived concept in CSP is the notion of *channel*. A channel is a collection of events of the form $c.n$: the prefix c is called the *channel name* and the collection of suffixes is called the *values* of the channel. When an event $c.n$ occurs it is said that *the value n is communicated on channel c*. When the value of a communication on a channel is determined by the environment (external choice) it is called an *input* and when it is determined by the internal state of the process (internal choice) it is called an *output*. It is convenient to write $c?n : N \rightarrow P(n)$ to describe behaviour over a range of allowed inputs instead of the longer $\Box\, n : N \bullet c.n \rightarrow P(n)$. Similarly the notation $c!n : N \rightarrow P(n)$ is used instead of $\sqcap\, n : N \bullet c.n \rightarrow P(n)$ to represent a range of outputs. Expressions of the form $c?n$ and $c!n$ do not represent events, the actual event is $c.n$ in both cases.

The interrupt process $P_1 \,\triangledown\, e \rightarrow P_2$ behaves as P_1 until the first occurrence of interrupt event e, then the control passes to P_2.

Recursion is used to given finite representations of non-terminating processes. The process expression

$$\mu\, P \bullet a?n : \mathbb{N} \rightarrow b!f(n) \rightarrow P$$

describes a process which repeatedly inputs a natural on channel a, calculates some function f of the input, and then outputs the result on channel b. CSP specifications are typically written as a sequence of simultaneous equations in a finite collection of process variables. Such a specification $\vec{X} \mathrel{\widehat{=}} \vec{F}(\vec{X})$ is implicitly taken to describe the solution to the vector recursion $\mu\, \vec{X} \bullet \vec{F}(\vec{X})$.

In general, the behaviour of a process at any point in time may be dependent on its internal state and this may conceivably take an infinite range of values. It is often not possible to provide a finite representation of a process without introducing some notation for representing this internal process state. The approach adopted by CSP is to allow a process definition to be parameterised by state variables. Thus a definition of the form

$$P_{n:N} \mathrel{\widehat{=}} Q(n)$$

represents a (possibly infinite) family of definitions, one for each possible value of n. There is no inherent notion of process state in CSP, but rather these annotations are a convenient way to provide a finite representation of an infinite family of process descriptions.

2.4.2 Timing primitives

To the standard CSP process primitives, Timed CSP adds two time specific primitives, the delay and the timeout.

A process which allows no communications for period t then terminates is written $\text{WAIT}\ t$. The process

$$\text{WAIT}\ t;\ P$$

is used to represent P delayed by time t.

The timeout construct passes control to an exception handler if no event has occurred in the primary process by some deadline. The process

$$(a \rightarrow P) \triangleright \{t\}\ Q$$

will try to perform $a \rightarrow P$, but will pass control to Q if the a event has not occurred by time t, as measured from the invocation of the process.

2.4.3 The timed collection example

The timed collection can be modeled as a process with two channels, *left* and *right* respectively. Elements are added to the collection through communications on the *left* channel and removed through communications on the *right* channel. The timing issues of the timed-collection can be described using (Timed CSP's) delay and timeout constructs.

The initial state is represented by the empty set.

$$TimedCollection \mathrel{\widehat{=}} TC_{\varnothing}.$$

When the first element joins the collection it is stamped with a timeout and the time taken to update the process state is represented by a delay.

$$TC_{\varnothing} \mathrel{\widehat{=}} left?e : X \rightarrow \text{WAIT}\ t_a;\ TC_{\{(t_o,e)\}}$$

When the collection is non-empty the process is ready to accept *left* or *right* events. The staleness stamps are updated with each communication and state update delays are introduced. In the event of no communication occurring before the oldest element becomes stale, all stale elements are purged.

$$\begin{array}{l} TC_{\{(t,a)\}\cup s} \mathrel{\widehat{=}} \\ \quad (left?e : X@t_i \rightarrow \text{WAIT}\ t_a;\ TC_{ps(t_i+t_a,\{(t,a)\}\cup s)\cup\{(t_o,e)\}} \ \Box \\ \quad right!a@t_i \rightarrow \text{WAIT}\ t_d;\ TC_{ps(t_i+t_d,s)}) \triangleright \{t\}\ \text{WAIT}\ t_p;\ TC_{ps(t+t_p,tl)} \end{array}$$

where $(t,a) = find_oldest(\{(t,a)\} \cup s)$.

$$\begin{array}{|l} [X] \\ \hline find_oldest : \mathbb{P}_1(\mathbb{T} \times X) \rightarrow (\mathbb{T} \times X) \\ \hline \forall s : \mathbb{P}_1(\mathbb{T} \times X) \bullet \exists (t,e) : s \bullet t = min(\operatorname{dom} s) \wedge find_oldest(s) = (t,e) \\ \hline \end{array}$$

2.4.4 Summary

For such an example Timed CSP is superior to Object-Z as a means of describing process control. Timed CSP also handles the timing issues of delays and timeouts simply and elegantly. The allowed sequences of events are clearly and concisely determined by the CSP code, there is no need to calculate preconditions nor is any other form of deep reasoning required to understand the ways in which the timed-collection may evolve. The Timed Object-Z approach results in a too complex model which over specifies this simple system, even though the timed collection example does not make use of the multi-threading and synchronisation capabilities of Timed CSP which are clearly well beyond the scope of Object-Z's atomic state transition semantics. On the

other hand, the syntactic treatment of internal state in the above is complex and unwielding, distracting strongly from the basically elegant treatment of the delay and timeout issues. Although, for example, Roscoe's CSP_M language [38] includes some powerful data modeling primitives, CSP still has no standard support for state modeling in the form of mathematical toolkits and libraries nor are there modular techniques for constructing and reasoning about complex internal state.

3 TCOZ

In many ways, Object-Z and Timed CSP complement each other in their capabilities. Object-Z has strong data and algorithm modeling capabilities. The Z mathematical toolkit is extended with object oriented structuring techniques. Timed CSP has strong process control modeling capabilities. The multi-threading and synchronisation primitives of CSP are extended with timing primitives. Moreover, both formalisms are already strongly influenced by the other in their areas of weakness. Object-Z supports a number of primitives which have been inspired by CSP notions such as external choice and synchronisation. CSP practitioners tend to make use of notation inspired by the Z mathematical toolkit in the specification of processes with internal state. This is not surprising given their joint associations in the Programming Research Group, Oxford. Another important connection is the well-known duality between the state transition behavioural model and the event based behavioural model [25] which makes it a simple matter to develop complementary semantics for the two languages.

Given these factors it is natural to consider the possibility of blending the two notations into a more complete approach to modeling real-time and/or concurrent systems. Fischer [19] and Smith [43] have independently suggested CSP-style semantics for Object-Z classes in which operation calls become CSP events. Operation names take on the role of CSP channels, with input and output parameters being passed down the operation channel as values. This view fits nicely with the Object-Z interpretation of operations being atomic, but is not well suited to considering multi-threading and real-time. Restricting operations to atomic events collapses the spatial and temporal aspects of operations, everything happens at a single point and instantaneously. Identifying channel names with operation names creates unnecessary tensions between the data and process views of objects and considerably reduces the potential for reuse of operation definitions.

The approach taken in the TCOZ notation is to identify operation schemas with terminating CSP processes that perform only state update events; to identify (active) classes with non-terminating CSP processes; and to allow arbitrary (channel-based) communications interfaces between objects.

The syntactic implications of this approach is that the basic structure of a TCOZ document is the same as for Object-Z. A document consists of a sequence of definitions, including type and constant definitions in the usual Z style. TCOZ varies from Object-Z in the structure of class definitions, which may include CSP channel and processes definitions. In fact, all operation definitions in TCOZ are considered to define CSP processes. The CSP view of an operation schema is that it describes all the sequences of update events which change the system state as required by the schema predicate. The exact nature and granularity of these update events is left undetermined in TCOZ (at least at the syntactic level), but by allowing an operation to consist of a number of events, it becomes feasible to specify its temporal properties when describing the operation. Another advantage of this approach is that by combining simple operations with CSP operators it becomes possible to represent true multi-threaded computation even at the operation level. The Fischer/Smith approach of identifying operation names with CSP channels is not followed, channels are given an independent, first class role. This allows the communications and control topology of a network of objects to be designed orthogonally to their class structure. The CSP channel mechanism is the only (dynamic) way to pass information between objects as the state of objects is encapsulated by hiding all update events.

3.1 Defining operations

The operation schema is the basic tool for describing state change in TCOZ. In order to allow treatment of timing issues in schema definitions, a distinguished identifier δ is introduced to represent the duration of the

state calculations performed by the operation. When δ does not appear in the definition of an operation, the default interpretation is that there be no bound on the duration of the operation, although individual specification documents may choose to adopt a different convention.

Although the schema is the basic tool, the true power of TCOZ comes from the ability to make use of Timed CSP primitives is describing the process aspects of an operation's behaviour. All operation definitions in TCOZ are in fact Timed CSP process definitions, with operation schema being given the syntactic status of terminating Timed CSP processes.

As an example, consider the specification of the *Add* operation of the timed-collection example. The actual state-change allowed by the operation schema remains unchanged from the Timed Object-Z version, but the timing characteristics of the operation are expressed by the condition $\delta = t_a$, rather than $now' - now = t_a$.

$$
\begin{array}{|l}
\hline
Add_0 \\
\Delta(mems) \\
e? : X \\
t_i? : \mathbb{T} \\
\hline
\delta = t_a \land mems' = ps(t_i? + t_a, mems) \cup \{(t_o, e?)\} \\
\hline
\end{array}
$$

Since TCOZ operations are identified with terminating CSP processes, it is natural to allow their definition in terms CSP primitives, such as event sequencing, as well as through the schema calculus. The novelty of the full TCOZ version of *Add* lies in the adoption of CSP primitives in its definition. Item inputs are communicated to the *Add* operation along a channel *left*.

$$Add \mathrel{\widehat=} [e : X;\ t_i : \mathbb{T}] \bullet left?e@t_i \rightarrow Add_0$$

This definition of *Add* says that after the parameter e has been input on channel *left* at time t_i, the state-calculation Add_0 is performed. Several aspects of TCOZ name-space conventions are raised by this definition.

Firstly, observe that the parameters e and t_i occur in Add_0 with Z-style input decorations and in *Add* without them. In TCOZ the convention is adopted that the true name of all parameters is the undecorated version. The ? and ! decorations are used solely in operation schemas, to distinguish between inputs and outputs. This is analogous to the convention of using primed and unprimed versions to indicate the final and initial values of a state attribute. In fact, all parameters are treated in the same way as state attributes, with the exception that state attributes are available in every name environment in a class definition.

This leads to the second observation that may be made of the *Add* definition. The $[e : X;\ t_i : \mathbb{T}] \bullet _$ construct is a local block definition in the state guard style (state guards are explained further below). The other forms of local blocks are the intentionial forms of both internal and external choice, which use the usual Z-style schema-text conventions. For example, $\bigsqcap n : \mathbb{N} \mid n < 5 \bullet c!n \rightarrow P$ or $\Box n : \mathbb{N} \mid n < 5 \bullet c?n \rightarrow P$. The state guard serves as an alternate form of external choice, so that the *Add* process is equivalent to $\Box e : X;\ t_i : \mathbb{T} \bullet left?e@t_i \rightarrow Add_0$.

In TCOZ, the local name space may be changed either by a local block definition as above or else by the occurance of an operation schema. An operation schema removes all its input parameters from scope and replaces them with its output parameters. The output parameters then become available for use in subsequent communication events or as inputs to subsequent operation schemas.

In the case of the *Delete* operation, the communication of the deleting element must precede the updating of the collection state and in fact is the enabling event for the operation. Since the name convention is that outputs are only available to the right of a schema, this behaviour cannot be described using an output parameter. Instead, the update operation is described as a simple state update which removes the oldest item (and any others that become stale).

$$
\begin{array}{|l}
\hline
Delete_0 \\
\Delta(mems) \\
t_i? : \mathbb{T} \\
\hline
mems \neq \varnothing \land \delta = t_d \land mems' = ps(t_i + t_d, mems \setminus \{(t, oldest)\}) \\
\hline
\end{array}
$$

The overall delete operation consists of this schema guarded by a communication on the *right* channel.

$$Delete \mathrel{\widehat=} [t_i : \mathbb{T} \mid mems \neq \varnothing] \bullet right!oldest@t_i \rightarrow Delete_0$$

The first part of the definition of *Delete* is a novel process control primitive known as a *state guard*.[5] The adoption of a state guard mechanism allows TCOZ to adopt a proper seperation between algorithm and process design issues. The sequencing of activities in an object is controlled explicitly through state guards rather than implicitly through the operation preconditions. In this way it becomes possible to reclaim the Z-style operation design and decomposition techniques abandoned by standard Object-Z.

Every process definition has (at least) an initial state which may be addressed using schema notation. This is the function of the first part of the expression defining *Delete*. It is a schema-based method of restricting the action of the process to initial states for which the collection is non-empty. For other states this process will *deadlock* or *block*, which is to say refuse any communication.

Note that the precondition requirement in the $Delete_0$ schema, though identical, could not achieve the desired restriction on the behaviour of *Delete*. Failure to satisfy a precondition when control is passed to an operation instead results in *divergence*, which is to say unspecified subsequent behaviour. $Delete_0$ places no restrictions at all on its behaviour when the initial queue is empty. The precondition is the state-based equivalent of process divergence and the guard is the state-based equivalent to process deadlock.

For every operation P (even those constructed using the process calculus) the collection of initial states for which the process will not diverge is called its *precondition* (written $\mathrm{pre}\ P$) and the collection of states for which it will not deadlock is called its *guard* (written $\mathrm{grd}\, P$).

3.2 Schemas and Processes

A schema expression describes a relationship on or between process state/s, whilst a process expression describes the overall behaviour or evolution of a process. The Z semantic model for operation schemas consists of sets of variable *bindings*, mappings from variable names to values. One semantic model for Timed CSP processes consists of sets of tuples consisting of a *trace* (a sequence of time stamped events), a *refusal* (a record of when events are refused by the process), and a *divergence* (a record of if and when the process diverged). The trace/refusal pair is called a *failure* and the overall model the failures/divergences model. The basic approach taken in the TCOZ semantics is to adopt the Timed CSP semantic model and to provide an interpretation of the Z semantic model in terms of failures and divergences, though two additions are required to make this possible. Firstly, a variable binding is added to represent the initial values of all the process attributes. Secondly, a new class of events, referred to as *update* events, is introduced to represent changes to the process attributes. The resulting model is called the state/failures/divergences model. The state of the process at any given time is the initial state updated by all of the updates that have occurred up to that time. If an event trace terminates (that is if a $\checkmark$ event occurs), then the state at the time of termination is called the *final* state.

The process model of an operation schema consists of all initial states and update traces (terminated with a $\checkmark$) such that the initial state and the final state satisfy the relation described by the schema. If no legal final state exists for a given initial state, the operation diverges immediately.

The process model for the state guard consists of replacing the trace part of every behaviour of the guarded process whose initial state does not satisfy the state guard with the empty trace. The empty event trace describes

[5] Although if-then style commands appear in several dialects of CSP, for example CSP_M [38], we believe that TCOZ is unique in adopting the state guard as a seperate primitive in the style of Morgan's version of the guarded command language [37].

the process being blocked by the failure of the state guard. In addition, divergence cannot occur if the state guard is not satisfied.

Since schema calculus operators cannot sensibly be applied to arbitrary CSP processes, it is necessary to strictly distinguish the schema calculus from the process calculus (see Appendix A). The two exceptions to this are the type-casting of operation schema expressions as terminating processes and of initial state schema expressions as state guards. In all other circumstance the schema and process calculi are separate and distinct. For example, if P and Q are operations schema expressions, the expression $a \rightarrow (P \land Q)$ is legal whilst the expression $(a \rightarrow P) \land Q$ is not. The full power of the schema calculus may be used to construct schema expressions, but once a schema expression has been cast into a process-like role it may no longer act in a schema-like role.

Some existing Object-Z schema calculus operators, such as $_ \Box _$, $_ \parallel _$, and $_ ; _$, have name-sakes with similar semantics in the CSP process calculus. The convention adopted in TCOZ is that the CSP operator is intended, only 'pure logic' schema calculus operators are allowed in TCOZ. This is justified by the superior algebraic properties of the CSP operators.

When operations are combined using the concurrency primitives $_ \parallel _$ and $_ \,|||\, _$, the designer is exposed to all the usual dangers of shared variable concurrency. The operation $OS_1 \parallel OS_2$, where OS_1 and OS_2 are operation schema, will synchronise on all state update events on variables in the respective delta-lists. Thus $OS_1 \parallel OS_2$ will have much the same process properties as $OS_1 \land OS_2$, with the exception that when the operations are inconsistent for a given initial state, the concurrent composition will deadlock while the logical composition with diverge. For example, consider

$$SQRT \mathrel{\widehat=} [x, x' : \mathbb{N} \mid x > 0 \land x'^2 = x]$$
$$HALF \mathrel{\widehat=} [x, x' : \mathbb{N} \mid x' * 2 = x].$$

The operations *SQRT* and *HALF* can agree only in the case where $x = 4$. When either of the operations is undefined (for example when $x = 2$ *SQRT* is undefined) $SQRT \parallel HALF$ will diverge. When both are defined but in disagreement (for example when $x = 16$) $SQRT \parallel HALF$ will deadlock at some unspecified time. The process $SQRT \land HALF$ is just $[x, x' : \mathbb{N} \mid x = 4 \land x' = 2]$ and will never deadlock. The concurrent composition $OS_1 \,|||\, OS_2$ is even less well behaved, every variable may be updated in any way allowed by either OS_1 or OS_2. Such a situation is likely to be very difficult to analyse. We strongly recommend that concurrent composition of operations be used sparingly, preferably only in cases where the operations have disjoint delta-lists. Shared data structures should only be utilised when properly protected by the object encapsulation mechanism.

3.3 Active and passive objects

The definition in a class of the distinguished process name MAIN indicates that the class is being defined as *active*. The MAIN process is used to determine the behaviour of objects of an active class after initialisation. Initialisation is treated in the usual way through the INIT schema. Active objects have their own thread of control and their mutable state attributes and operation definitions are fully encapsulated (update events are hidden). Distinct objects, even of the same class, share no data and can experience no shared variable interference. Other objects can neither reference an active object's state attributes nor invoke any of its local operations. Only local constants, such as the object identity attribute *self*, may be accessed by other classes. All dynamic interactions with an active object must take place through the CSP channel communication mechanism. Active objects are considered to have the syntactic properties of process identifiers and may be composed using CSP operators.

The MAIN operation is optional in a class definition. If a class is defined without a MAIN process it is called a *passive* class. Passive objects are controlled by other objects in a system and their state and operations are fully available to the controlling object (unless explicitly hidden). The appearance of MAIN clearly distinguishes the definition of active objects and passive objects in a system.

Returning to the timed-collection example, the existence of environmental obligations and the need to purge stale elements means that the timed-collection class must have its own thread of control. Assuming that the class operations are defined as in Section 3.1, the timed-collection behaviour is defined by a MAIN process

similar to the Timed CSP version presented in Section 2.4.3.

$$\textsc{Main} \mathrel{\widehat=} \mu\, TC \bullet [mems = \varnothing] \bullet Add;\ TC \ \Box \\ [mems \neq \varnothing] \bullet ((Add \ \Box\ Delete) \triangleright \{t\}\ Purge);\ TC$$

The most striking difference lies in the use of the operation schemas to subsume the role of the complex annotations present in the Timed CSP version. This represents a clearer and more structured presentation of the basic control logic. A second difference lies in the use of the state guard construct to distinguish between the empty and non-empty behaviours of the timed-collection, thus saving the need to define separate empty and non-empty processes.

3.4 Communication channels

The class state-schema convention is extended to allow the declaration of communication channels. If c is to be used as a communication channel by any of the operations of a class, then it must be declared in the state schema to be of type **chan**. Channels are type heterogeneous and may carry communications of any type.

Aside This convention has been criticised by a number of reviewers on the grounds that it violates the strong-typing of Z and denies the reader possibly instructive information. We believe that the ability to send many forms of data over a channel plays a vital role in reducing the complexity of class interfaces and can be justified conceptually by the observation that it is more instructive to group logically related communications (such as those pertaining to a particular protocol) than to group communications with identical type but logically unrelated function. In any case, all data communicated on channels remains strongly-typed.

Contrary to the conventions adopted for internal state variables, channels are viewed as global rather then as encapsulated entities. This is an essential consequence of their role as communications interfaces *between* objects. In the situation of multiple instances of objects of the same class in a system, those objects will all share the same channel. For example, if $\mathcal{O}$ is a sequence of objects of class C with channel c, then in the process

$$|||\, i : \operatorname{dom} \mathcal{O} \bullet \mathcal{O}(i)$$

each of the objects $\mathcal{O}(i)$ communicates with the environment by sharing channel c with every other object. In the general case there is no way for the environment to know which of the objects it is communicating with when using channel c. If it is necessary to know which object the environment is communicating with, the object identity attribute *self* [15] can be included in the communication, for example

$$c.(self, message).$$

(This technique is used frequently in the lift case study.) The introduction of channels to TCOZ reduces the need to reference other classes in class definitions, thereby further enhancing the modularity of system specifications.

Consider once again the timed-collection example. Using the TCOZ conventions, the class state can be significantly simplified from the Timed Object-Z version. The sole remaining primary class attribute is the actual collection itself, none of the timing attributes are required. In addition to the list of *mems*, the state schema must declare channels *left* and *right*. These channels serve much the same role as the corresponding environment variables in the Timed Object-Z version, but here that role is better defined in terms of the CSP process model. The secondary attributes *oldest* and t remain useful in simplifying the operation definitions and are retained.

$$\begin{array}{|l} mems : \mathbb{F}\,(\mathbb{T} \times X) \\ left, right : \textbf{chan} \\ \Delta \\ (t, oldest) : \mathbb{T} \times X \\ \hline mems \neq \varnothing \Rightarrow (t, oldest) \in mems \wedge t = min \operatorname{dom} mems \\ \hline \end{array}$$

3.5 The timed-collection

Bringing together the various aspects of the TCOZ timed-collection introduced above, we are able to present the entire class definition.

TimedCollection[X]

$mems : \mathbb{P}(\mathbb{T} \times X)$
$left, right :$ **chan**
Δ
$(t, oldest) : \mathbb{T} \times X$

$mems \neq \varnothing \Rightarrow (t, oldest) \in mems \land t = min \operatorname{dom} mems$

INIT

$mems = \varnothing$

Add$_0$

$\Delta(mems)$
$e? : X$
$t_i? : \mathbb{T}$

$\delta = t_a \land mems' = ps(t_i? + t_a, mems) \cup \{(t_o, e?)\}$

Delete$_0$

$\Delta(mems)$
$t_i : \mathbb{T}$

$mems \neq \varnothing \land \delta = t_d$
$mems' = ps(t_i + t_d, (mems \setminus \{(t, oldest)\})$

$Add \mathrel{\widehat=} [e : X;\ t_i : \mathbb{T}] \bullet left?e@t_i \rightarrow Add_0$
$Delete \mathrel{\widehat=} [t_i : \mathbb{T} \mid mems \neq \varnothing] \bullet right!oldest@t_i \rightarrow Delete_0$

Purge

$\Delta(mems)$

$\delta = t_p \land mems' = ps(t + t_p, mems)$

MAIN $\mathrel{\widehat=} \mu TC \bullet [mems = \varnothing] \bullet Add;\ TC \ \Box$
$[mems \neq \varnothing] \bullet ((Add \ \Box \ Delete) \triangleright \{t\}\ Purge);\ TC$

This specification represents a more concise, flexible, and scalable treatment of both process and state than is possible in either Object-Z or Timed CSP. The structure of the process' internal state and communications interfaces are prominently documented. The structured schema based approach to describing state transitions, supported as it is by the full power of the Z toolkit and the schema calculus, is better able to handle large and complex process state than the essentially *ad hoc* state annotation conventions of CSP. Making use of the Timed CSP process definition conventions removes the need to consider process control matters in operation schemas. There is a clear separation of process control and algorithmic matters which simplifies the description of both.

3.6 Composing classes

Inheritance is a mechanism for incremental specification, whereby new classes may be derived from one or more existing classes.

Essentially, all definitions are pooled with the following provisions. Inherited type and constant definitions and those declared in the derived class are merged. The state and initialisation schemas of inherited classes and those declared in the derived class are conjoined. Operation schemas with the same name are also conjoined.

TCOZ extends Object-Z with two new class constructs, channels '**chan**' and main behaviour 'MAIN'. Channels are treated as normal state constant attributes, therefore, they are pooled into the derived classes. Channel renaming is the same as state attribute renaming. Since new classes will generally have new behaviours, the MAIN class definition is never inherited. As for all other class definitions, a class extension must include a MAIN definition if the class is to be active. The rules for (active and passive) class inheritance are:

- A new active class may be derived from an existing active class by defining a new MAIN process.

- A new active class may be derived from an existing passive class by defining a MAIN process.
- A new passive class can also be derived from an existing active class by not defining a new MAIN process.
- A new passive class can be derived from an existing passive class by not defining a MAIN process.

A composite object which contains active objects is also an active object (a MAIN definition must appear in the composite object class). Active objects are responsible for their own intialisation, so a composite object will often not require an explicit INIT schema. An the example of this is the Buffered-Consumer/producer class in Section 3.7. Since it has no passive internal state it requires no explicit initialisation. Another result of the encapsulation of local state in active objects is the inability of a composite class to refer to the local state attributes and operations of a component object. Only class constants such as the object identity *self* may be accessed.

3.7 A multi-threaded example

The timed-collection example made no use of the multi-threaded capabilities of the TCOZ notation. In this section, a specification of a standard buffered consumer/producer process is presented as a demonstration these aspects of TCOZ.

A simple consumer/producer process accepts inputs of type L on its *left* channel, calculates some function

$$f : L \rightarrow R$$

of the inputs, and outputs the result on its *right* channel.

ConProd

$left, right$: **chan**

Process

$i? : L$

$o! : R$

$\delta = t_p \land o! = f(i?)$

MAIN $\widehat{=}$ $\mu C \bullet [i : L] \bullet left?i \rightarrow Process;\ right!o \rightarrow C$

To ensure that all inputs are accepted and outputs are received, the process is buffered left and right with timed-collection processes. The buffered process consists of a left buffer, a right buffer, and an internal consumer/producer, as depicted in Figure 1.

Correct hookup of the timed-collection buffers to the consumer/producer is achieved by *renaming* the various internal channels. The renaming convention is the same as for Object-Z and Timed CSP, that is $P[a/b]$ is P with all occurrences of b replaced by a. Intermediate channels *ml* and *mr* are introduced as the internal interfaces to the left and right buffers respectively. In order to retain a consumer/producer like interface, the *right* channel of the left buffer is renamed to *ml*, the *left* channel of the right buffer renamed to *mr* and the *left* and *right* channels of the internal consumer/producer are renamed to *ml* and *mr* respectively.

The internal interfaces are protected from environmental interference by *hiding* them. The hiding notation is the same as for both Object-Z and Timed CSP, that is $(P \setminus c)$ is P with c protected from external influence. In the case where P is process-like and c is a channel this has the important result of freeing communications on c from the requirement of synchronising with the environment. Thus communications on *mr* and *ml* occur as soon as the local processes are ready and cannot be blocked by any other entity.

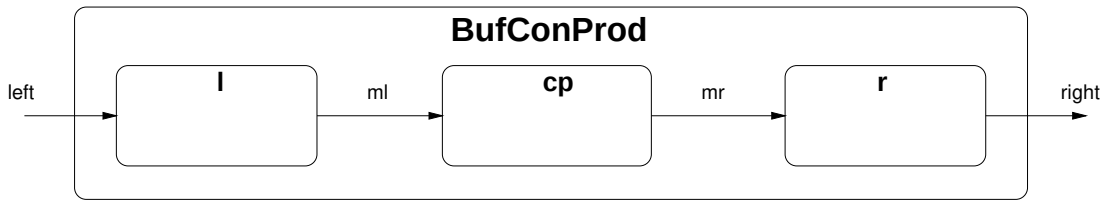

Figure 1: Buffered consumer producer

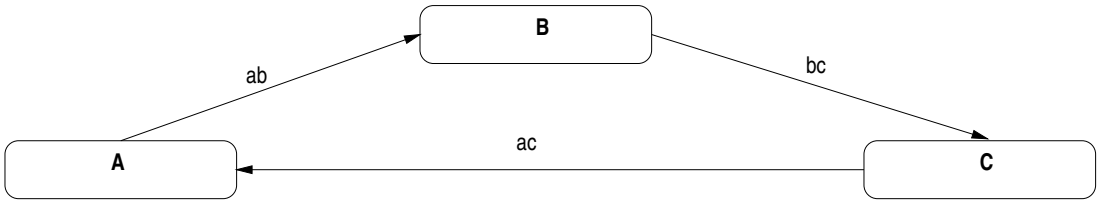

Figure 2: Two dimensional communication topology

BufConProd

$l : TimedCollection[L][ml/right]$
$r : TimedCollection[R][mr/left]$
$cp : ConProd[ml/left, mr/right]$

$\text{MAIN} \mathrel{\widehat=} (l \,|[\, ml \,]|\, cp \,|[\, mr \,]|\, r \setminus ml, mr)$

The *BufConProd* class definition allows true multithreading of the two buffers and the consumer/producer. For example, the left buffer may be accepting a new item at the same time as the consumer/producer is processing another item, at the same time as the right buffer is releasing yet another item to the environment. This concurrency is coerced into smooth co-operation through the requirement for synchronisation between the processes when communication occurs on the internal channels *ml* and *mr*.

The addition of these CSP process structuring features represents a significant advance in the scope and size of systems which may be addressed by the Object-Z approach.

3.8 Complex network topologies

The syntactic structure of the CSP synchronisation operator is convenient only in the case of pipe-line like communication topologies. Expressing more complex communication topologies generally results in unacceptably complicated expressions. For example, consider the communication topology shown in Figure 2, processes A and B communicate privately through the channel ab, processes A and C communicate privately through the channel ac, and processes B and C communicate privately through the channel bc. One CSP expression for such a network communication system is

$$(A[bc'/bc] \,|[\, ab, ac \,]|\, (B[ac'/ac] \,|[\, bc \,]|\, C[ab'/ab]) \setminus ab, ac, bc)[ab, ac, bc/ab', ac', bc'].$$

The hiding and renaming is necessary in order to cover cases such as C being able to communicate on channel ab.

The above expression not only suffers from syntactic clutter, but also serves to obscure the inherently simple network topology described so elegantly by Figure 2. One reason for this is that it artificially suggests a dominant role for process A. Equivalent, but superficially very different, formulations of the network could be used

to assign this "dominant" role to either B or C. We believe that network topologies can be better described by adopting a notation inspired by the graph-based approach embodying in Figure 2.

A *network topology abstraction* is is an expression of the form

$$(\Big|\Big| v_1, v_2, v_3 \ldots \bullet v_1 \overset{ch_{12}}{\longleftrightarrow} v_2;\ v_2 \overset{ch_{23}}{\longleftrightarrow} v_3;\ v_3 \overset{ch_{13}}{\longleftrightarrow} v_1;\ \ldots).$$

The variables $v_1, v_2, v_3, \ldots$ are called the *formal network parameters* and the network connections

$$v_1 \overset{ch_{12}}{\longleftrightarrow} v_2;\ v_2 \overset{ch_{23}}{\longleftrightarrow} v_3;\ v_3 \overset{ch_{13}}{\longleftrightarrow} v_1;\ \ldots$$

are called the *formal network topology*.

The formal network parameters are abstract names that represent the process that form the nodes of the network topology. It is necessary to use abstract names to allow separate incarnations of processes with identical definition. The formal network topology describes a finite graph in which each edge or *network connection* generally represents a private channel connection between network nodes. Multiple connections between processes and connections between multiple processes over a single channel are represented by multiple connection expressions. In this case the channel becomes a party line between the various participating processes.

A network topology abstraction describes a function, which we call a *network constructor*, that builds a process network using its process arguments as nodes. The network constructor associated with a given network abstraction is built by processing the formal parameters sequentially. At each stage there is a list of remaining formal parameters. Suppose v is the next remaining formal parameter. The local channels on which v is not synchronising are decorated to avoid unwanted synchronisations. The local channels on which it is synchronising with any of the remaining parameters are used to create a synchronisation with the remaining network. Then the rest of the formal parameters are treated one at a time to construct an expression for that remaining network. Finally all the local channels are hidden and any decorations removed.

The above network topology abstraction thus describes the following network constructor.

$$\begin{array}{l}
\lambda\, v_1, v_2, v_3, \ldots \bullet \\
\quad (v_1[ch'_{23}/ch_{23}, \ldots]\,|[\,ch_{12}, ch_{13}, \ldots\,]| \\
\qquad (v_2[ch'_{13}/ch_{13}, \ldots]\,|[\,ch_{23}, \ldots\,]| \\
\qquad\quad (v_3[ch'_{12}/ch_{12}, \ldots]\,|[\,\ldots\,]| \\
\qquad\qquad \vdots \\
\qquad\quad) \\
\qquad) \\
\quad \setminus\, ch_{12}, ch_{13}, ch_{23}, \ldots)[ch_{12}/ch'_{12}, ch_{13}/ch'_{13}, ch_{23}/ch'_{23}, \ldots]
\end{array}$$

The system in Figure 2 can be described by applying a suitable network topology abstraction to the processes A, B, and C.

$$(\Big|\Big| v_1, v_2, v_3 \bullet v_1 \overset{ab}{\longleftrightarrow} v_2;\ v_2 \overset{bc}{\longleftrightarrow} v_3;\ v_3 \overset{ac}{\longleftrightarrow} v_1)(A, B, C)$$

The processes A, B, C are the *actual network parameters*. When the actual network parameters are all process names, the syntactic conventions are relaxed to allow the formal network topology to act in the guise of a process operator. For example, the network topology of Figure 2 may be described by the lax usage

$$\Big|\Big|(A \overset{ab}{\longleftrightarrow} B;\ B \overset{bc}{\longleftrightarrow} C;\ C \overset{ca}{\longleftrightarrow} A).$$

Such usage is considered acceptable since the names representing the actual parameter can serve the dual purpose of also identifying the formal parameters. Other forms of lax usage allow network connections with common nodes to be run together, for example

$$\Big|\Big|(A \overset{ab}{\longleftrightarrow} B \overset{bc}{\longleftrightarrow} C \overset{ca}{\longleftrightarrow} A),$$

and multiple channels above the arrow, for example

$$\Big|\Big| (A \xleftrightarrow{ab_1, ab_2} B).$$

As a further example of the use of the network topology convention, note that the behaviour of the *BufConProd* object can also be expressed by

$$\textsc{Main} \mathrel{\widehat{=}} \Big|\Big| (l \xleftrightarrow{ml} cp \xleftrightarrow{mr} r).$$

The network topology convention is particular useful for describing the interactions between floors (requests), lift-controller and elevators in the lift system.

We believe this TCOZ network topology convention to be a good candidate to be included in the basic CSP notation itself.

4 The lift case study

The multi-lift system is a standard example used to demonstrate the expressive power of various specification techniques in modeling concurrent reactive systems.[6] People are familiar with the user requirement of a lift system, so they can concentrate on the modeling notations. However, the lift case study is not a trivial example because of the complexity caused by inherent concurrent interaction in the system [45]. We chose the specification of the lift system as the TCOZ case study also because both CSP and Object-Z have been applied to the lift system allowing a comparison to be drawn. The CSP 'lift' model [41] describes the sequences of events for the lift system well, however, it struggles to capture the data aspects of the lift system. Furthermore, the CSP model has a flat and in places awkward structure and the communications interfaces between the lift system components (i.e. floor-buttons and lifts) are not clearly documented. The Object-Z 'lift' model [13] demonstrates the power of modeling the state change of the lift system in a structured way. However, it is complicated by the need to represent process state as data and it uses a complex, centralised control-model because of the Object-Z's single thread semantics. Neither the CSP or Object-Z model addresses the real-time issues for the lift system.

Our goals for the TCOZ specification of the lift system include:

- a model that captures both the data structures and the behaviour of the lift system;
- a true multi-threaded specification that captures the concurrent, reactive, and real-time aspects of the lift system; and
- a component based incremental style specification that is extendible and reusable.

This case study aims to demonstrate the TCOZ approach for modeling timed reactive systems. Other detailed (lift specific) issues, such as the efficiency of lift scheduling, are not considered in this paper.

4.1 System overview

A lift system for a building consists of multiple lifts each providing transport between the various floors of the building as dictated by the pressing of a range of service-request buttons. Inside each lift there is a panel of buttons, one for requesting travel to each of the building's floors. The panel buttons of any lift must be in one-to-one correspondence with the floor numbers. In general there are two service-request buttons on each floor, for upward and downward travel respectively, though on the first floor and the top floor there is only one button.

[6] Many formal notations have been applied to the lift system. For example, the UNITY model [6, 8], Raddle [20], and Constraint Nets [46].

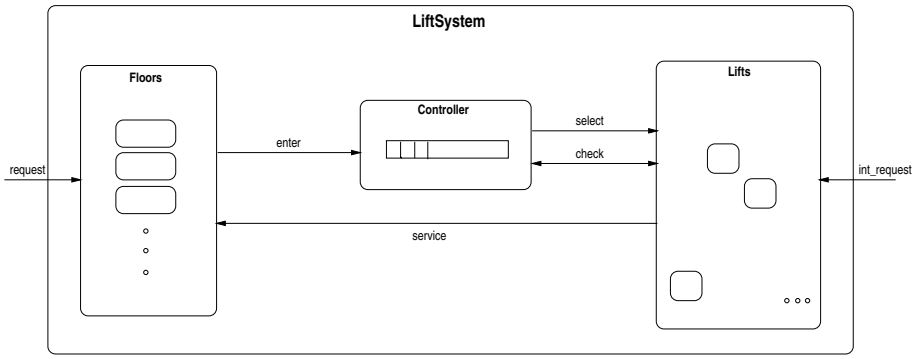

Figure 3: Lift system communication diagram

Any service-request button can be pushed at any time. Once pushed the button is said to be *on* and it remains on until the requested service is provided. Pressing an internal button requests the lift to visit the corresponding floor. Pressing an external button requests a lift to visit the floor with the desired direction of travel. The lift controller has a queue which stores all current (external) floor requests. When a request arrives from a floor, the system will put the request at the end of the external request queue. When a lift becomes available, the controller will assign the first request of the queue to the lift for service. When visiting a floor, a lift door operates in the order of open-door, wait, then close-door. This normal process can be interrupted by a customer crossing the door as detected by some sensors.

Furthermore, the following timing properties must be captured in the model:

- lift travel time between two consequent floors is a constant, however there is a constant time delay for acceleration and braking;
- without interrupts, the lift door should be kept at the 'open' state for a fixed time period before closing.

4.2 Specification structure

The lift specification is developed in a bottom-up manner, beginning with models of basic (passive) component objects, such as 'buttons', which are then used to develop more complex (active) component objects such as 'floors (requests)', 'lifts' and a 'controller'. Then the lift system is modeled as an active composite object which composes the component objects with their interactions (through channels).

Figure 3 illustrates the communication interfaces between lift system components. There are three major components, the service-request panels on each floor, the lifts themselves, and the central controller which mediates service requests from the floors. External requests received by a floor on the *request* channel cause the floor's corresponding service button to be lit and the request is communicated to the controller on the *enter* channel. The button remains lit until a confirmation is received on the *service* channel. Requests received by the controller on the *enter* channel are enqueued and sent to idle lifts on the *select* channel on a first-in-first-out basis. Whenever a lift receives an internal request on the *int_request* channel, the corresponding button is lit and the requested entered into its itinerary. Whilst the lift has local requests pending it services them in strict order according to its current movement direction, reversing direction at the extreme floors. The behaviour of the active lifts is monitored on the *check* channel and if a request can be serviced en route it is dispatched to the lift in question and dequeued. Once the lift becomes idle, it may accept an external request on the *select* channel, move to the requested floor, and send a confirmation to the floor panel on the *service* channel.

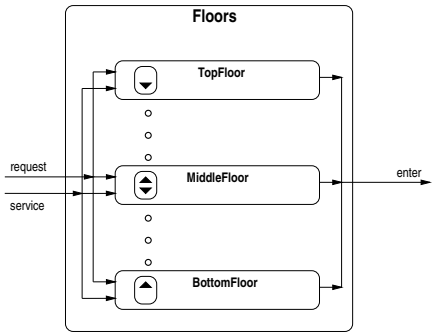

Figure 4: External service request panels.

4.3 Buttons

A basic component of the lift system is the button panels on the floors or inside the lifts. Buttons have a common behaviour; they can be pushed 'On' by people and turned 'Off' by the system.

$$ButtonStatus ::= On \mid Off$$

Buttons are modeled using a simple passive class which records their current state and provides operations for turning them on and off.

Button

$state : ButtonStatus$

INIT

$state = Off$

TurnOn

$\Delta(state)$

$state' = On$

TurnOff

$\Delta(state)$

$state' = Off$

4.4 The building

Our model of the building concentrates on the behaviour of the service-request panels on each floor, as depicted in Figure 4, other floor related issues are elided.

4.4.1 Floor panels

Floors may be divided into two classes, those from which it is possible to travel upward and those from which is possible to travel downward. The *TopFloor* is a floor from which only downward travel is possible,

TopFloor

$num : \mathbb{N}$

downbutton : $Button_{\copyright}$
request, *enter*, *service* : **chan**

INIT
downbutton.INIT

$PressDown \mathrel{\widehat=} request?(num, Down) \rightarrow$
$\quad [downbutton.state = Off] \bullet downbutton.TurnOn;\ enter!(num, Down) \rightarrow \text{SKIP} \ \Box$
$\quad [downbutton.state = On] \bullet \text{SKIP}$
$DownOff \mathrel{\widehat=} service?(num, Down) \rightarrow downbutton.TurnOff$
$\text{MAIN} \mathrel{\widehat=} \mu T \bullet (PressDown \ \Box \ DownOff);\ T$

and the *BottomFloor* is a floor from which only upward travel is possible.

BottomFloor

$num : \mathbb{N}$

upbutton : $Button_{\copyright}$
request, *enter*, *service* : **chan**

INIT
upbutton.INIT

$PressUp \mathrel{\widehat=} request?(num, Up) \rightarrow$
$\quad [upbutton.state = Off] \bullet upbutton.TurnOn;\ enter!(num, Up) \rightarrow \text{SKIP} \ \Box$
$\quad [upbutton.state = On] \bullet \text{SKIP}$
$UpOff \mathrel{\widehat=} service?(num, Up) \rightarrow upbutton.TurnOff$
$\text{MAIN} \mathrel{\widehat=} \mu B \bullet (PressUp \ \Box \ UpOff);\ B$

A *MiddleFloor* is a floor from which both upward and downward travel is possible. Object-Z's class inheritance features are used to allow both upward and downward travel for the *MiddleFloor* class.

MiddleFloor

TopFloor, *BottomFloor*

$downbutton \neq upbutton$

$\text{MAIN} \mathrel{\widehat=} \mu M \bullet (PressDown \ \Box \ DownOff \ \Box \ PressUp \ \Box \ UpOff);\ M$

Inheritance provides a subclassing mechanism for specification reuse (not a true subtyping mechanism). When a new active class is derived from an existing active class, the objects of the new class often have different behaviour to the objects of the existing class. As a consequence, the MAIN process must always be redefined explicitly.

The subscript '$_{\copyright}$' (object containment[7] [11]) indicates that the button objects are contained in their corresponding floor object.

As a floor can be either top-floor or bottom-floor or middle-floor, the general type of a floor is defined as a class-union [9].

$$Floor \mathrel{\widehat=} TopFloor \cup BottomFloor \cup MiddleFloor$$

4.4.2 The building

The building is modeled as an aggregate of active floor objects. Individual floors do not communicate with each other, but rather with the central controller and with the lifts. Thus the MAIN processes of the individual floor objects are combined using asynchronous composition, $_ \,|||\, _$.

[7] Object containment ensures that no object directly or indirectly contains itself; and no object is directly contained in two distinct objects. For a detailed discussion see [11].

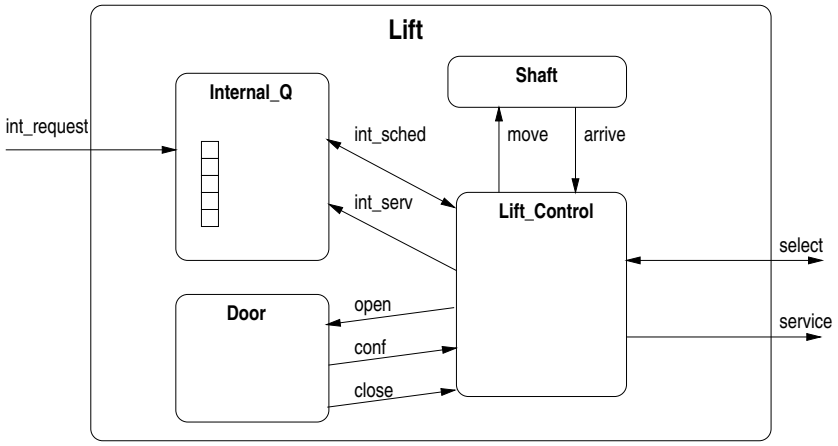

Figure 5: Internal lift communication diagram

Building

$height : \mathbb{N}$

$floors : \operatorname{seq} Floor_{\copyright}$

$height = \#floors$
$\forall i : \operatorname{dom} floors \bullet floors(i).num = i$
$floors(1) \in BottomFloor \land floors(height) \in TopFloor$
$\forall i : (2 \mathinner{..} (height - 1)) \bullet floors(i) \in MiddleFloor$

$\textsc{Main} \mathrel{\widehat=} |||\, fl : \operatorname{dom} floors \bullet floors(fl)$

4.5 Lifts

A lift consists of four parts as depicted in Figure 5, a door for allowing access to and from the lift, a shaft for transporting the lift, an internal queue for determining the lift itinerary, and a controller for coordinating the behaviour of the other components. This division structures the specification in such a way as to limit the complexity of the individual components and to highlight the potential for concurrency.

4.5.1 Lift door control

The lift door controller is treated as a separate class so as to ensure a clear description of its timing and safety properties. Under this limited aim, the class may be described entirely within the Timed CSP idiom.

The controller interfaces on a channel *servo* with a servomechanism that activates the door to open or close and on a channel *sensor* to determine when the door is open, closed, or blocked from closing. The messages that may be set on these channels are then

$$DoorMess ::= ToOpen \mid Opened \mid ToClose \mid Closed \mid Interrupt.$$

The timing property of the door is that once it is open, it must remain open for time period

$t_o : \mathbb{T}$ [open time]

before closing. The safety property is that if the closing of the door is blocked (as indicated by receipt of an *Interrupt* message) the door must be reopened.

The door cycle is initiated by receipt of an *open* signal from the lift controller and completed by sending a *close* signal. As soon as the door is open, a *conf* signal must be sent to the lift controller so as to indicate fulfillment of a service request.

Door

$$open, conf, close, servo, sensor : \textbf{chan}$$

$$OpenDoor \mathrel{\widehat=} servo!(self, ToOpen) \rightarrow sensor?(self, Opened) \rightarrow \text{SKIP}$$
$$CloseDoor \mathrel{\widehat=} servo!(self, ToClose) \rightarrow sensor?(self, Closed) \rightarrow \text{SKIP}$$
$$CycleDoor \mathrel{\widehat=} OpenDoor;\ conf \rightarrow$$
$$\quad (\mu\, CD \bullet \text{WAIT}\, t_o;\ CloseDoor \triangledown (sensor?(self, Interrupt) \rightarrow OpenDoor;\ CD))$$
$$\text{MAIN} \mathrel{\widehat=} \mu\, D \bullet open \rightarrow CycleDoor;\ close \rightarrow D$$

In the *CycleDoor* process we make use of the CSP interrupt primitive which allows an exception handling behaviour to be triggered by the occurrence of an unusual event. The normal door cycle follows the order of open-door, wait t_o, then close-door. This normal cycle can be interrupted and re-started by the event which detects a message *Interrupt* from the channel *sensor*. This message acts as an interrupt on the normal cycle and control is returned to the start of the normal cycle.

4.5.2 Moving the lift

The essentially analog nature of the movement of the lift presents something of a modeling problem. The technique used throughout this specification has been to abstract real-world interactions as CSP events (eg the *request* and *int_request*) channels, but the movement betweens floors is by its nature a time-consuming process. We thus adopt the common technique of delimiting the period of movement by start and finish events. The start event of the movement process is a communication to the lift-shaft apparatus of the number of floors to be moved. The finish event is a communication from the lift-shaft apparatus that the lift has arrived at the destination floor. The timing properties of lift movement are described by two time constants.

$$\bar{t} : \mathbb{T} \qquad \text{[maximum time to move one floor Up or Down]}$$
$$delay : \mathbb{T} \qquad \text{[acceleration and braking delay]}$$

The maximum time to pass from one floor to another is $\bar{t}$ for each floor travelled plus a delay of *delay* caused by initial acceleration and final braking of the lift.

Shaft

$$move, arrive : \textbf{chan}$$

$$\text{MAIN} \mathrel{\widehat=} \mu\, S \bullet [n \neq 0] \bullet move?n \overset{0<\delta\leq|n|*\bar{t}+delay}{\rightarrow};\ arrive \rightarrow S$$

Such event based models are highly abstract and perhaps are less satisfying when applied to the complex process of moving the lift, than when applied to the more event-like process of pressing a button. However, the channel based interfaces of TCOZ processes mean that such models must be used. In their favour it must be pointed out that from the point of view of the lift controller the matters of essential interest are precisely when the movement commences and when it finishes.

4.5.3 Lift itinerary

The itinerary of the lift is determined by the requests made by passengers using the internal floor-request panel and those dispatched from the central control. For the purposes of limiting design complexity a separate class is defined to determine the lift's internal service itinerary.

This internal queue makes use of a *panel* of buttons to communicate with passengers and to maintain a record of the floor-requests pending. A dependent (secondary) variable *irs* records the set of destinations that have been selected by the passengers at any given time. This set may in turn be split into those destinations above a given floor, *ups*, and those below, *dns*. Floor services are requested on the *int_request* channel and confirmed on the *int_serv* channel. Scheduling services are requested by passing the current floor *fl*? and movement direction *md*?

$$MoveDirection ::= Up \mid Down$$

on the *int_sched* channel.

The operations of the controller are turning the panel buttons on and off, in response to service requests and confirmations, and determining the next destination for the lift itinerary. The next destination is the first requested destination in the current direction of movement, reversing at either extremity of movement. A scheduling request is only serviced if there are pending floor-service requests.

Internal_Q

$height : \mathbb{N}$

$panel : \operatorname{seq} Button_{©}$
$int_request, int_sched, int_serv : \textbf{chan}$
Δ
$irs : \mathbb{F}\ \mathbb{N}$
$ups, dns : \mathbb{N} \pfun \mathbb{F}\ \mathbb{N}$

$height = \#panel$
$irs = \{f : \operatorname{dom} panel \mid panel(f).state = On\}$
$ups = \lambda fl : \operatorname{dom} panel \bullet \{n : irs \mid n \geq fl\}$
$dns = \lambda fl : \operatorname{dom} panel \bullet \{n : irs \mid n \leq fl\}$

NextUp

$fl?, dest! : \operatorname{dom} panel$

$(ups(fl?) \neq \varnothing \land dest! = min\ ups(fl?))$
$\lor$
$(ups(fl?) = \varnothing \land dns(fl?) \neq \varnothing \land$
$dest! = max\ dns(fl?))$

NextDown

$fl?, dest! : \operatorname{dom} panel$

$(dns(fl?) \neq \varnothing \land dest! = max\ dns(fl?))$
$\lor$
$(dns(fl?) = \varnothing \land ups(fl?) \neq \varnothing \land$
$dest! = min\ ups(fl?))$

$Next \mathrel{\widehat=}$
$[md? : MoveDirection \mid md? = Up] \land NextUp \lor$
$[md? : MoveDirection \mid md? = Down] \land NextDown$

$\textsc{Main} \mathrel{\widehat=} \mu IQ \bullet$
$[fl : \operatorname{dom} panel] \bullet int_request?(self, fl) \rightarrow panel(fl).TurnOn;\ IQ\ \Box$
$[fl : \operatorname{dom} panel] \bullet int_serv?fl \rightarrow panel(fl).TurnOff;\ IQ\ \Box$
$[fl : \operatorname{dom} panel;\ md? : MoveDirection \mid irs \neq \varnothing] \bullet$
$int_sched?(fl, md) \rightarrow Next;\ int_sched!dest \rightarrow IQ$

4.5.4 Lift controller

The lift controller keeps record of the current floor and movement direction and provides the interface between the lift environment and the other lift components. The lift controller exhibits of three modes of behaviour.

The lift begins at rest awaiting either a passenger destination request or a dispatch from the central controller. Any passengers inside the lift are given a period of time

$$| \; t_p : \mathbb{T} \qquad \text{[wait for passenger input]}$$

to make an internal request before the lift accepts any external requests.

The controller determines that an internal request is pending through the willingness of the internal queue to perform a scheduling transaction. Once the queue has indicated the next destination the central controller is checked to see if there is an external request from an intermediate floor. If so the liftservices the external request first. If not it services the internal request: the new direction is calculated and set; the lift is moved to the new floor and the door is opened; the internal queue and the central controller are notified; and then once the door is closed control is returned to the rest mode.

If, after waiting t_p, an external request becomes available before any internal request: the new move direction is set; the lift is moved to the new floor and the door is opened; the floor service-request panel is notified; and then, once the door is closed, control is returned to the rest mode.

LiftControl

$fl : \mathbb{N}$
$md : MoveDirection$
$move, arrive :$ **chan** [shaft]
$open, conf, close :$ **chan** [door]
$int_sched, int_serv :$ **chan** [internal Q]
$select, check, service :$ **chan** [external channels]

SetFloor

$\Delta(fl)$
$dest? : \mathbb{N}$

$fl' = dest?$

SetDir

$\Delta(md)$
$dir? : MoveDirection$

$md' = dir?$

CalDir

$dest? : \mathbb{N}$
$dir! : MoveDiection$

$dest? = fl \land dir! = md \lor$
$dest? > fl \land dir! = Up \lor$
$dest? < fl \land dir! = Down$

$$Move \mathrel{\widehat=} [dest : \mathbb{N} \mid dest \neq fl] \bullet move!(dest - fl) \rightarrow arrive \rightarrow open \rightarrow conf \rightarrow SetFloor \; \Box \; [dest : \mathbb{N} \mid dest = fl] \bullet open \rightarrow conf \rightarrow \textsc{Skip}$$

$$Internal \mathrel{\widehat=} CalDir;\ SetDir;\ Move;\ int_serv!fl \rightarrow close \rightarrow \textsc{Skip}$$

$$Get_Internal \mathrel{\widehat=} int_sched!(fl, md) \rightarrow [dest : \mathbb{N}] \bullet int_sched?dest \rightarrow check!(fl, dest, md) \rightarrow check \rightarrow Internal \; \Box \; [dest : \mathbb{N}] \bullet check?dest \rightarrow External$$

$$External \mathrel{\widehat=} Move;\ service!(fl, md) \rightarrow close \rightarrow \textsc{Skip}$$

$$Get_External \mathrel{\widehat=} [dest : \mathbb{N};\ dir : MoveDirection] \bullet select?(dest, dir) \rightarrow SetDir;\ External$$

$$\textsc{Main} \mathrel{\widehat=} \mu LC \bullet Get_Internal \; \Box \; (\textsc{Wait}\, t_p;\ Get_External);\ LC$$

4.5.5 The bank of lifts

Each lift consists of a door, a shaft, and a controller.

Lift

$iq : Internal_Q_{\copyright}$
$lc : LiftControl_{\copyright}$
$s : Shaft_{\copyright}$
$d : Door_{\copyright}$

$\textsc{Main} \mathrel{\widehat=} \left|\left|\right| (lc \xleftrightarrow{move,arrive} s;\ lc \xleftrightarrow{open,close,conf} d;\ lc \xleftrightarrow{int_sched,int_serv} iq)$

The collection of all the lifts in the system is modeled as an aggregate of the individual lifts acting autonomously.

Lifts

$lifts : \mathbb{P}\, Lift_{\copyright}$

$\textsc{Main} \mathrel{\widehat=} \left|\left|\left|\right|\right| l : lifts$

4.6 The central controller

The responsibility of the central controller is to dispatch floor requests to idle lifts. It consists of a request queue with channels that connect the floors and the lifts. The network topology is described graphically in Figure 3.

Controller

$reqQ : \mathrm{seq}(\mathbb{N} \times MoveDirection)$
$enter, select, check, service :$ **chan**
Δ
$req : \mathbb{N} \leftrightarrow MoveDirection$
$ups, dns : \mathbb{N} \times \mathbb{N} \to \mathbb{P}\, \mathbb{N}$

$req = \mathrm{ran}\, reqQ$
$\forall f_1, f_2 : \mathbb{N} \bullet$
$\quad ups(f_1,f_2) = \{f : \mathbb{N} \mid f \mathrel{\underline{req}} Up \land f_1 \le f \le f2\}$
$\quad dns(f_1,f_2) = \{f : \mathbb{N} \mid f \mathrel{\underline{req}} Down \land f_1 \le f \le f2\}$

Init

$reqQ = \langle\,\rangle$

Join

$\Delta(reqQ)$
$req? : (\mathbb{N} \times MoveDirection)$ **on** $enter$

$reqQ' = reqQ \frown \langle req? \rangle$

Remove

$\Delta(reqQ)$
$(dest?, md?) : \mathbb{N} \times MoveDirection$

$reqQ' = squash(reqQ \rhd \{(dest?, md?)\})$

Find

$fl?, d? : \mathbb{N}$
$md? : MoveDirection$
$dest! : \mathbb{N}$
$suc! : \mathbb{B}$

$md? = Up \land suc! = (ups(fl?, d?) \neq \varnothing) \land$
$suc! \Rightarrow dest! = \min ups(fl?, d?)$
$\lor$
$md? = Down \land suc! = (dns(fl?, d?) \neq \varnothing) \land$
$suc! \Rightarrow dest! = \min dns(fl?, d?)$

$Dispatch \mathrel{\widehat=}$
$\quad [(dest, md) : \mathbb{N} \times MoveDirection \mid reqQ \neq \langle\,\rangle \land (dest, md) = head\, reqQ] \bullet$
$\quad\quad select!(dest, md) \to Remove$

$$
\begin{array}{l}
\mathit{CheckServ} \mathrel{\widehat=} \\
\quad [\mathit{fl}, d : \mathbb{N};\ \mathit{md} : \mathit{MoveDirection}] \bullet \mathit{check}?(\mathit{fl}, d, \mathit{md}) \rightarrow \mathit{Find}; \\
\quad\quad [\mathit{suc}] \bullet \mathit{check}!\mathit{dest} \rightarrow \mathit{Remove} \ \Box \\
\quad\quad [\neg\ \mathit{suc}] \bullet \mathit{check} \rightarrow \textsc{Skip} \\
\textsc{Main} \mathrel{\widehat=} \mu\, C \bullet (\mathit{Join} \ \Box\ \mathit{Dispatch} \ \Box\ \mathit{CheckServ});\ C
\end{array}
$$

The controller receives requests from the floors and enters them into the *reqQ* queue (*Join* operation). In the ordinary case these requests are dispatched in first-in-first-out manner (*Dispatch* operation) as idle lifts become available, but if in the course of servicing internal requests a lift can visit a floor whilst moving in the floor required direction, the request is removed from the queue (*CheckServ* operation).

4.7 The lift system

The lift system consists of the floors of the building, the bank of lifts and the central controller. The number of floor-service buttons in each lift must be exactly the number of floors in the building.

LiftSystem

$$
\begin{array}{l}
bd : \mathit{Building} \\
ls : \mathit{Lifts} \\
\mathit{contr} : \mathit{Controller} \\
\hline
\forall\, l : ls.\mathit{lifts} \bullet l.\mathit{iq}.\mathit{height} = bd.\mathit{height} \\
\hline
\textsc{Main} \mathrel{\widehat=} \Big|\Big| (bd \overset{\mathit{enter}}{\longleftrightarrow} \mathit{contr} \overset{\mathit{select},\mathit{check}}{\longleftrightarrow} ls \overset{\mathit{service}}{\longleftrightarrow} bd)
\end{array}
$$

The lift system behaviour MAIN describes the communication channels between the independent concurrently executing system components: the lifts, the floors, and the controller. The floors communicate service requests to the controller through the *enter* channel, the controller dispatches these requests to the lifts through the *select* and *check* channels, and the lifts indicate successful servicing of requests to the floors through the *service* channel.

4.8 Discussion

The application of TCOZ to the multi-lift system has been convincingly successful, despite the relatively modest real-time aspects of the specification. The powerful combination of object-oriented structuring, data modeling, and process modeling features available in TCOZ has allowed the clean presentation of an ambitiously detailed (when compared to versions described in other languages) treatment of the multi-lift system. This strongly supports our claims that TCOZ represents a highly scalable and reusable method for describing real-time and concurrent systems. As an example, consider modifying the specification so as to allow 'opportunistic' servicing of requests entered after a lift starts moving. The modular nature of the specification immediately draws attention to the *Shaft* class which controls the movement of the lift between floors. The interface of this class needs to be expanded to include events to indicate the lift is approaching the next floor and to stop the lift at the next floor. Armed with these additional controls the *LiftControl* class can easily be modified to react opportunistically to new service requests.

One surprise in the development of the TCOZ lift was the degree to which the process model idiom came to dominate. The starting point of the development had been a data oriented standard Object-Z specification [13]. Rather than being a simple matter of adding real-time and concurrent features to the existing specification it soon became clear that much of its 'data' was in fact being used to represent 'process' behaviour which could

be more elegantly represented using the CSP process modeling features of TCOZ. The final approach adopted was to model the system primarily as a network of communicating processes and to make use of Object-Z's data modeling features to simplify and structure the specification by abstracting away from algorithmic specifics and reusing common data components such as buttons.

A perceived weakness of the TCOZ approach was identified in its handling of the interface between TCOZ processes and the real world. Although the abstraction of button presses as communications on external channels is reasonably acceptable, the inability to describe 'continuously' changing aspects of the system such as 'lift position' is particularly disturbing. Whilst modeling the moving process by start and finish events provides an adequate interface to the TCOZ specification it goes no way at all toward ensuring that lift shaft satisfies our informal intuitions as to its behaviour. The specification is thus strictly speaking not of a system for moving people between floors of a building, but rather simply a description of a method of controlling such a system in a satisfactory manner.

One advantage of choosing the lift system to exercise the TCOZ language is the availability of existing lift specifications in both CSP [41] and Object-Z [13]. Apart from the ability of TCOZ to describe timing aspects not addressed by either of these specifications, such as the correct behaviour of the door opening cycle, TCOZ represents an improvement in expressibility, modularity, and reusability over both existing specifications.

The CSP 'lift' model [41] is similar in spirit to the TCOZ lift presented here, except that it does not consider the door cycle nor the movement of the lift. However, in contrast to the highly modular approach of the TCOZ specification, the CSP version is forced to adopt a quite flat structure because the only structuring facility available is the process definition. The channel interface declarations and the network topology operator provide valuable information to the reader of the TCOZ specification regarding the source and destination of communications, which is not available to the reader of the CSP specification. Although some attempt is made to structure the CSP specification document through the use of section headings to distinguish system components, the essentially global nature of all process definitions and communication channels makes it difficult to comprehend individual process definitions without extensive reference to the rest of the specification. The standard CSP synchronisation operator ($_ \parallel _$) is a particular point of weakness as it offers no visual feedback as to the interface between network components. The use of the network topology operator (and associated diagrams) is a particular strength of the TCOZ specification.

The weaknesses in CSP's treatment of data and algorithms not only adds syntactic clutter to the CSP lift, but also appears to influence the structure of the specification. In particular, the difficulty in abstracting various calculations away from the specific direction of travel results in a 'split' specification, with many features being repeated for both the 'up' and 'down' directions of travel. For example, the specification of a single lift [41, p 26] includes subprocesses $Lift(f, up)$ and $Lift(f, down)$ which differ primarily in the method of calculating the next floor on the itinerary. In contrast the TCOZ version abstracts this calculation into the *Next* operation of the *Internal_Q* class allowing the description of the lift's gross behaviour to be independent of the direction of travel, even despite using a more complex method of determining the itinerary. Although a specification of this form is possible in CSP, the lack of strong, modular data modeling facilities acts as a strong disincentive to this form of abstraction.

The Object-Z 'lift' model [13] provided a structured and reusable model and was able to describe the sequence of state changes of the overall lift system. However, Object-Z's process semantics forces the specification to include extra data for describing process state and all system components to be viewed passively (except the lift system class) which leads to a complex, centralised control-model. The lift components, such as the shaft and the door, were abstractively modeled as internal state components (rather than component objects) of the lift. The system class *LiftSystem* [13, pp 146,147] became very complex (and lengthy) because the global ordering of synchronisations between lifts, floor requests and request queue must be explicitly determined. In the TCOZ model only the local ordering of these events need be specified, the global ordering is implicit in the CSP event synchronisation model. In addition, the TCOZ model gives the freedom of viewing the lift system components as active entities allowing a more natural, modular, and reusable description of the lift system.

5 Related work

The basis for the successful blending of the Timed CSP and Object-Z notations is the duality between state transition semantics and event semantics. This has long been recognised [25, 5] and has undoubtedly helped shape the development of Object-Z's behavioural semantics. Perhaps the most mature formalism based on this duality is Butler's [5] blending of CSP with Back's Action Systems [3]. An important lesson from this work is the need to distinguish strongly between the notions of guard and precondition in the state-transition view. The failure to do so in Object-Z has made it impossible to reconcile the usual Z algorithmic refinement techniques with the default behavioural semantics, greatly reducing the value of these techniques in Object-Z. The adoption of a distinct notion of state guards in TCOZ makes possible a full blending of Z-style algorithm and CSP-style process refinement.

The notion of blending the untimed CSP and Object-Z has been proposed independently by Fischer [19] and Smith [43]. Both take the approach of identifying the notion of channel with that of operation and operation invocations with atomic communications of both inputs and outputs. The latter prevents the modeling of timing and concurrency at the operation level and complicates the CSP semantics through the mixing of elements of external (inputs) and internal (outputs) choice in a single event. The former is undesirable from a theoretical standpoint because it confuses communications interfaces which are essentially process related attributes with algorithmic structures which are essentially data related attributes. An object's communications interface should be determined by high-level considerations of the overall system structure, whilst the operational interface should be determined by consideration of the internal data structures. The purpose of the class envelope is to resolve such tensions locally, not to propagate them up and down the design hierarchy. The practical consequences of the identification of channel and operation is the promotion of both high degrees of coupling between classes and unnatural class structures. Neither formalism makes a complete distinction between preconditions, guards, and operations and consequently refinement issues are complicated in both. Smith adopts a semantics which is unable to model process divergence and as a consequence must identify preconditions with guards, making process and algorithmic refinement incompatible. The semantics adopted by Fischer does allow a distinction between guards and preconditions, but guards are tightly coupled with operations so that the same operation may not be used in differing circumstances as is the *Remove* operation in the *Controller* class. Moreover, a convention is introduced whereby when an operation guard is not explicitly defined the precondition is used by default, thus complicating both the understanding of the process behaviour and the refinement of the operation. Issues, such as real-time and the distinction between active and passive objects, are not addressed by either formalism.

More generally, the need for specification notations capable of addressing both data/algorithmic issues and process control issues is widely recognised. Several notations now exist aimed at bridging this divide. These fall essentially into two classes, those that adopt a process-algebra/event-based style (LOTOS [29], ESTEREL [4], RAISE [22]) and those that adopt a transition system style (UNITY [6], Action Systems [3], TLA [31]). We consider two real-time specific languages, E-LOTOS [28] and ASTRAL [7], as being representative of their respective classes and having similar aims to TCOZ.

The LOTOS specification language [29] is very similar in approach to TCOZ, blending CSP-like process primitives with a declarative-style data-specification language. E-LOTOS [28] is a recently developed real-time extension to LOTOS. The process and real-time primitives of E-LOTOS are influenced by Timed-CSP, therefore these aspects of E-LOTOS are similar to TCOZ. The major differences lie in the data-modeling and structuring aspects of the two formalisms. The data modeling language of E-LOTOS is an algebraic/functional hybrid, whilst TCOZ is model based. The module construct of E-LOTOS is similar to the class construct of TCOZ in that it can encapsulate states and operations. Modules can be reused via the *imports* mechanism which is similar to the class inheritance. However E-LOTOS modules cannot be instantiated as a type, while TCOZ classes can. Therefore, the notions of object and composition of objects (aggregation) are missing from E-LOTOS. In TCOZ this adds another dimension of potential for reuse of specifications. E-LOTOS's subtyping is a simple record-type extension mechanism which is less powerful than the TCOZ's polymorphic typing (inheritance hierarchy and class union).

ASTRAL is developed based on the ideas of RT-ASLAN [2] and TRIO [23]. ASTRAL also has a module construct that encapsulates the variables and transactions. The ASTRAL module has the two dimensions of reuse, importing and type instantiation. Therefore the ASTRAL module is very close to the TCOZ class construct except that the names of instances of a module are modeled explicitly, while the object identity is implicitly included in the TCOZ class semantics (a class is a collection of object identities). In ASTRAL, object composition is generally modeled by the (programming language flavoured) list type construct, *array ...of ...*. In TCOZ, object composition may have various (mathematics flavoured) abstract forms, i.e. collections '$\mathbb{P}\ \ldots$' and list '$\mathrm{seq}\ \ldots$'. The timing aspects are similar to the TLA approach [1] and the Timed Object-Z approach [10]. That is, a global clock *NOW* is introduced and system environments are modeled explicitly. Therefore sequential real-time systems can be captured well by ASTRAL, while truly concurrent active systems are difficult to describe in ASTRAL.

6 Conclusions and further work

Timed CSP and Object-Z complement each other not only in their expressive capabilities, but also in their underlying semantics. In addition, the object oriented flavour of Object-Z provides an ideal foundation for promoting modularity and separation of concerns in system design. The combination of the two, TCOZ, treats data and algorithmic concerns in the Object-Z style and treats process control, timing, and communication concerns in the CSP style. The notion of active and passive objects are clearly distinguished in the TCOZ model.

This powerful modeling combination, TCOZ, has been successfully applied to a comprehensive case study on specifying a real-time multi-lift system. In comparison to the CSP model [41] and the standard Object-Z model [13] of the lift system, the TCOZ model not only describes the complex system state and behaviour within a cleaner and less coupled structure, but also captured the true concurrent real-time interactions between various system components of the lift system. The lift case study also provides feedback to the development of TCOZ. For example, the development of the modeling notation for complex network topologies is motivated by the lift case study. A particular weakness of the language has been identified in its ability interface with 'real-world' aspects of a system. This clearly limits the applicability of the notation to the software aspects of a system. Future work will be directed toward improving its capabilities in this direction so as to integrate it into a more holistic approach to real-time and embedded systems design. One promising approach is to enhance TCOZ with features of the Timed Refinement Calculus [32, 36] which allow convincing descriptions of continuously varying real-world observables.

TCOZ preserves in large part both the syntax and semantics of the individual notations and hence can potentially benefit from the large body of experience developed in the use of and tool support for the individual notations and their parents. These benefits might include: the full application of the Z schema calculus and associated structured design techniques, which are severely restricted in scope in Object-Z; Z verification and refinement techniques, again restricted in Object-Z; the application of Timed CSP process equivalence and refinement texhniques, especially the Timed CSP model hierarchy for moving between timed and untimed models; the extension of existing Z, Object-Z, and CSP tools and model checkers. For example, we are currently planning a project to construct a parser/type-checker for TCOZ based on the Object-Z parser of Johnson [30].

A separate paper details the blended state/event process model which forms the basis for the TCOZ semantics [33]. Additional planned work includes developing refinement rules for the TCOZ specification language based on existing Z and CSP refinement systems. Schneider has described a system for capturing and verifying abstract temporal requirements of Timed CSP processes [39], it is hoped that this might also form a valuable addition to the TCOZ notation.

Acknowledgements

We would like to thank John Colton, Ian Hayes and Keith Gallagher for many useful comments. This work is supported in part by the DSTO/CSIRO Fellowship programme.

References

[1] M. Abadi and L. Lamport. An Old-Fashioned Recipe for Real Time. In J.W. de Bakker, , C. Huizing, W.-P. de Roever, and G. Rozenberg, editors, *Proc. REX Workshop on Real-Time: Theory in Practice*, volume 600 of *LNCS*, pages 1–27. Springer-Verlag, 1991.

[2] B. Auernheimer and R.A. Kemmerer. Rt-aslan: A specification language for real-time systems. *IEEE Trans. Software Eng.*, 12(9), September 1986.

[3] R. J. R. Back and J. von Wright. Refinement calculus, part II: Parallel and reactive programs. In J. W. de Bakker, W. P. de Roever, and G. Rozenberg, editors, *Stepwise Refinement of Distributed Systems: Models, Formalism, Correctness*, volume 430 of *Lecture Notes in Computer Science*, pages 42–66. Springer Verlag, 1990.

[4] G. Berry and G. Gonthier. The ESTEREL synchronous programming language: design, semantics, implementation. *Science of Computer Programming*, 19(2):87–152, November 1992.

[5] M. J. Butler. *A CSP Approach to Action Systems*. PhD thesis, Wolfson College, Oxford University, Michaelmas Term 1992.

[6] K. M. Chandy and J. Misra. *Parallel Program Design: A Foundation*. Addison Wesley, 1988.

[7] A. Coen-Porisini, C. Ghezzi, and R. Kemmerer. Specification of Realtime Systems Using ASTRAL. Technical Report 96-30, Computer Science Department, University of California, Santa Barbara, US, January 1997.

[8] H. C. Cunningham, V.R.Shan, and S. Shen. Devising a formal specification for an elevator controller. TR 94-10, Department of Computer and Information Science, University of Mississippi, USA, 1994.

[9] J.S. Dong. Living with free type and class union. In *The 1995 Asia-Pacific Software Engineering Conference (APSEC'95)*, pages 304–312. IEEE Computer Society Press, December 1995.

[10] J.S. Dong, J. Colton, and L. Zucconi. A formal object approach to real-time specification. In *the 3rd Asia-Pacific Software Engineering Conference (APSEC'96)*, Seoul, Korea, December 1996. IEEE Computer Society Press.

[11] J.S. Dong and R. Duke. The geometry of object containment. *Object-Oriented Systems*, 2(1):41–63, Chapman & Hall, March 1995.

[12] J.S. Dong, G. Rose, and R. Duke. The role of secondary attributes in formal object modelling. In Alex Stoyenko, editor, *The First IEEE International Conference on Engineering Complex Computer Systems (ICECCS'95)*, pages 31–38, Ft. Lauderdale, USA, November 1995. IEEE Computer Society Press.

[13] J.S. Dong, L. Zucconi, and R. Duke. Specifying parallel and distributed systems in Object-Z. In G. Agha and S. Russo, editors, *The 2nd International Workshop on Software Engineering for Parallel and Distributed Systems*, pages 140–149, Boston, Massachusetts, 1997. IEEE Computer Society Press.

[14] R. Duke, P. King, G. Rose, and G. Smith. The Object-Z specification language: Version 1. Technical Report 91-1, Software Verification Research Centre, Australia, 1991.

[15] R. Duke and G. Rose. Modelling object identity. In *Proc. 16th Australian Comput. Sci. Conf. (ACSC-16)*, pages 93–100, February 1993.

[16] R. Duke and G. Rose. *Formal Object Oriented Specification*. Academic Press, 1998. forthcoming.

[17] R. Duke, G. Rose, and G. Smith. Object-Z: a specification language advocated for the description of standards. *Computer Standards and Interfaces*, 17:511–533, 1995.

[18] C. Fidge, P. Kearney, and M. Utting. A formal method for building concurrent real-time software. *IEEE Software*, 14(2), 1997.

[19] C. Fischer. CSP-OZ: A combination of Object-Z and CSP. In H. Bowmann and J. Derrick, editors, *Formal Methods for Open Object-Based Distributed Systems (FMOODS '97)*, volume 2, pages 423–438. Chapman & Hall, 1997.

[20] I. R. Forman. Design by decomposition of multiparty interactions in Raddle87. In *The 5th IEEE International Workshop on Software Specification and Design (IWSSD'89)*, pages 2–10. IEEE Computer Society Press, 1989.

[21] A. J. Galloway and W. J. Stoddart. An operational semantics for ZCCS. In M. Hinchey and S. Liu, editors, *the IEEE International Conference on Formal Engineering Methods (ICFEM'97)*, pages 272–282, Hiroshima, Japan, November 1997. IEEE Computer Society Press.

[22] C. George, P. Haff, K. Havelund, A. E. Haxthausen, R. Milne, C. Bendix Nielson, S. Prehn, and K. R. Wagner. *The Raise Specification Language*. Prentice Hall, New York, 1992.

[23] C. Ghezzi, D. Mandrioli, and A. Morzenti. Trio: A logic language for executable specifications of real-time system. *J. Systems and Software*, June 1990.

[24] I. J. Hayes and B. P. Mahony. Using units of measurement in formal specifications. *Formal Aspects of Computing*, 7(3), 1995.

[25] Jifeng He. Process simulation and refinement. *Formal Aspects of Computing*, 1(3):229–241, 1989.

[26] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall International, 1985.

[27] International Organization for Standardization, Geneva. *Units of measurement: handbook on international standards for units of measurement*, 1979.

[28] ISO. SC21/WG7 Working Draft on Enhancements to LOTOS, ISO Working Group 7. December 1997.

[29] ISO 8807. *LOTOS – A formal description technique based on the temporal ordering of observational behaviour*, 1989.

[30] W. Johnston. A type checker for Object-Z. Technical report 96-24, Software Verification Research Centre, School of Information Technology, The University of Queensland, Brisbane 4072. Australia, July 1996.

[31] L. Lamport. The temporal logic of actions. *ACM Transactions on Programming Languages and Systems*, 16:872–923, 1994.

[32] B. P. Mahony. Networks of predicate transformers. Technical Report 95-05, Software Verification Research Centre, Department of Computer Science, The University of Queensland, St. Lucia, QLD 4072, Australia, February 1995.

[33] B. P. Mahony and J.S. Dong. Blending Object-Z and Timed CSP: The semantics of TCOZ. Technical Report 97-24, Mathematical and Information Sciences, Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia, 1997.

[34] B. P. Mahony and J.S. Dong. Blending Object-Z and Timed CSP: An introduction to TCOZ. In *The 20th International Conference on Software Engineering (ICSE'98)*. IEEE Computer Society Press, April 1998. (accepted).

[35] B. P. Mahony and J.S. Dong. Network topology and a case-study in TCOZ. In *ZUM'98 The* 11^{th} *International Conference of Z Users*, volume 1493 of *Lecture Notes in Computer Science*. Springer-Verlag, September 1998. (accepted).

[36] B. P. Mahony and I. J. Hayes. A case-study in timed refinement: A mine pump. *IEEE Transactions on Software Engineering*, 18(9):817–826, 1992.

[37] C. C. Morgan. *Programming from Specifications*. Prentice Hall, second edition, 1994.

[38] A. W. Roscoe. *Theory and Practice of Concurrency*. International Series in Computer Science. Prentice-Hall, 1997.

[39] S. Schneider. *Correctness and Communication in Real-Time Systems*. PhD thesis, Oxford University Computing Laboratory, Programming Research Group, 1990. Available as Technical Monograph PRG-84.

[40] S. Schneider and J. Davies. A brief history of Timed CSP. *Theoretical Computer Science*, 138, 1995.

[41] M. D. Schwartz and N. M. Delisle. Specifying a lift control system with CSP. In *The 4th IEEE International Workshop on Software Specification and Design (IWSSD'87)*, pages 21–27, Monterey, California, April 1987. IEEE Computer Society Press.

[42] G. Smith. A fully abstract semantics of classes for Object-Z. *Formal Aspects of Computing*, 7(3):289–313, 1995.

[43] G. Smith. A semantic integration of Object-Z and CSP for the specification of concurrent systems. In *Proceedings of FME'97: Industrial Benefit of Formal Methods*, Graz, Austria, September 1997. Springer-Verlag.

[44] K. Taguchi and K. Araki. The State-Based CCS Semantics for Concurrent Z Specification. In M. Hinchey and S. Liu, editors, *the IEEE International Conference on Formal Engineering Methods (ICFEM'97)*, pages 283–292, Hiroshima, Japan, November 1997. IEEE Computer Society Press.

[45] J.C.P. Woodcock, S. King, and I.H. Sorensen. Mathematics for specification and design: The problem with lifts. In *The 4th IEEE International Workshop on Software Specification and Design (IWSSD'87)*, pages 265–268. IEEE Computer Society Press, 1987.

[46] Y Zhang and A. K. Mackworth. Design and analysis of embedded real-time systems: An elevator case study. Technical Report 93-04, Computer Science Department, University of British Columbia, 1993.

A TCOZ Concrete Syntax

The TCOZ syntax is based on the Object-Z syntax [14, 16].

A.1 Notation

The syntax is described in an extended BNF with the following metasymbols:

$::=$	produces
$\mid$	alternative
$[x]$	optional x
$\{x\}$	zero or more x's
$\{x\}_1$	one or more x's

Nonterminals ending in *List* have productions according to:

$$xList ::= x\{,x\}$$

Nonterminal names are typically compound and abbreviated, each part commencing with an upper-case letter, e.g. *OpnExpDef*. Abbreviations are listed below.

Terminal symbols are shown directly as they appear in TCOZ. Metasymbols are larger than terminal symbols of the same shape as seen by:

$::=$ ::= , $\mid$ |, $[\,]$ [ ] and $\{\}$ { }.

A.2 Abbreviations

The following abbreviations are used in the productions listed in Section A.3.

Abbrev	Abbreviation	Gen	Generic
Ax	Axiomatic	Invar	Invariant
Bool	Boolean	Opn	Operation
Chan	Channel	Pred	Predicate
Dec	Declaration	Prim	Primary
Def	Definition	Sec	Secondary
Des	Designator	Sep	Separator
Exp	Expression	Var	Variable

A.3 Productions

The order of presentation of productions is top down with definitions for nonterminals appearing after their last application, except for recursive definitions.

$$Specification ::= Def\{Sep\,Def\}$$

$$Def ::= ClassDef \mid NonClassDef$$

$$ClassDef ::= \begin{array}{|l} ClassHeading \\ \hline [Visibility] \\ [Inheritance] \\ [LocalDefs] \\ [StateSchema] \\ [InitialSchema] \\ [Operations] \\ \hline \end{array}$$

$$ClassHeading ::= ClassName\,[GenFormals]$$

$$Visibility ::= \upharpoonright (NameList)$$

$NameList ::= Name \{, Name\}$

$Inheritance ::= ClassDes \{Sep\ ClassDes\}$

$LocalDefs ::= NonClassDef \{Sep\ NonClassDef\}$

$NonClassDef ::= GivenTypeDef \mid FreeTypeDef \mid AbbrevDef \mid AxDef \mid GenDef$

$GivenTypeDef ::= [NameList]$

$FreeTypeDef ::= Name ::= Branch \{ \mid Branch\}$

$Branch ::= Name [\langle\!\langle Exp \rangle\!\rangle]$

$AbbrevDef ::= Name [GenFormals] \mathrel{\widehat{=}} Exp$

$AxDef ::=$

$$\begin{array}{|l} Decls \\ \hline [Pred] \\ \hline \end{array}$$

$GenDef ::=$

$$\begin{array}{|l} [GenFormals] \\ \hline Decls \\ \hline [Pred] \\ \hline \end{array}$$

$GenFormals ::= [NameList]$

$StateSchema ::=$

$$\begin{array}{|l} [ChanDecs] \\ [PrimVarDecs] \\ [\Delta SecVarDecs] \\ \hline [ClassInvar] \\ \hline \end{array}$$

$ChanDecs ::= ChanDec \{; ChanDec\}$

$ChanDec ::= NameList : \mathbf{chan}$

$PrimVarDecs ::= Decs$

$SecVarDecs ::= Decs$

$ClassInvar ::= Pred$

$InitialSchema ::=$

$$\begin{array}{|l} \textsc{Init} \\ \hline Pred \\ \hline \end{array}$$

$Operations ::= OpnDef \{Sep\ OpnDef\}$

$OpnDef ::= OpnSchemaDef \mid OpnExpDef$

$OpnSchemaDef ::=$

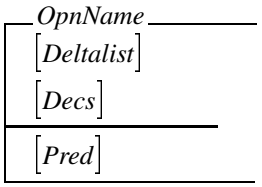

$OpnExpDef ::= OpnName \mathrel{\widehat{=}} OpnExp$

$OpnExp ::= \sqcap SchemaText \bullet OpnExp$

$\mid \Box SchemaText \bullet OpnExp$

$\mid ||| SchemaText \bullet OpnExp$

$\mid NetworkTopology$

$\mid [SchemaText] \bullet OpnExp$

$\mid OpnExp;\ OpnExp$ L

$\mid OpnExp \Box OpnExp$ L

$\mid OpnExp \sqcap OpnExp$ L

$\mid OpnExp \mid\mid\mid OpnExp$ L

$\mid OpnExp \mid [Exp] \mid OpnExp$ L

$\mid OpnExp \parallel OpnExp$ L

$\mid OpnExp \triangleright \{Exp\}\ OpnExp$ L

$\mid OpnExp \swarrow \{Event\}\ OpnExp$ L

$\mid OpnExp \triangledown\ OpnExp$ L

$\mid OpnExp \bullet \text{DEADLINE}\ Exp$

$\mid OpnExp \bullet \text{WAITUNTIL}\ Exp$

$\mid \text{WAIT}\ Exp$

$\mid Event[@Exp] \rightarrow OpnExp$

$\mid \mu\, OpnName \bullet OpnExp$

$\mid OpnExp\ Renaming$

$\mid OpnExp \setminus (NameList)$

$\mid OpnSchemaExp$

$\mid (OpnExp)$

$\mid [Exp\ .]\ OpnName$

$\mid \bot \mid \textsc{Stop} \mid \textsc{Skip}$

$OpnSchemaExp ::= \bigwedge SchemaText \bullet OpnExp$

$\mid \bigvee SchemaText \bullet OpnExp$

$\mid OpnExp \wedge OpnExp$ L

$\mid OpnExp \vee OpnExp$ L

$\mid [\ [Deltalist]\ [Decs]\ [\ \mid Pred]\ ]$

$\mid [Exp\ .]\ OpnName$

$\mid (OpnExp)$

$NetworkTopology ::= \|(Connection \{;\ Connection\})$

$Connection ::= NameList \xleftarrow{NameList} NameList$

$Event ::= Name[.Exp]$

Productions are in equal-precedence groups (separated by $----$) and the precedence of groups increases down the page. The 'L/R' indication determines left or right associativity of binary operators.

$OpnName ::= Name \mid \textsc{Main}$

$Deltalist ::= \Delta(NameList)$

$SchemaText ::= Decs[\ \mid Pred]$

$Decs ::= Dec \{;\ Dec\}$

$Dec ::= NameList : Exp$

$Pred ::= \forall SchemaText \bullet Pred$

$\mid \exists SchemaText \bullet Pred$

$\mid \exists_1 SchemaText \bullet Pred$

$\mid \mathrm{let}\ LetDefs \bullet Pred$

$\mid Pred \Leftrightarrow Pred$ L

$\mid Pred \Rightarrow Pred$ R

$\mid Pred \vee Pred$ L

$\mid Pred \wedge Pred$ L

$$
\begin{array}{l}
\quad --------- \\
\mid \neg\ Pred \\
\mid Name\ .\ \textsc{Init} \\
\mid \mathsf{true} \mid \mathsf{false} \\
\mid BoolExp \\
\mid (Pred)
\end{array}
$$

The nonterminal *Deltalist* cannot be written *DeltaList* as the construct represented does not conform to the *List* convention.

Iff ($\Leftrightarrow$), disjunction ($\vee$) and conjunction ($\wedge$) are associative. 'R' indicates right-to-left association.

$$
\begin{array}{l}
BoolExp\ ::=\ Exp \\
Exp\ ::=\ \mu\, SchemaText \left[\ \bullet\ Exp\right] \\
\quad \mid \lambda\, SchemaText \bullet Exp \\
\quad \mid \mathrm{let}\ LetDefs \bullet Exp \\
\quad \mid \mathrm{if}\ Pred\ \mathrm{then}\ Exp\ \mathrm{else}\ Exp \\
\quad\quad --------- \\
\quad \mid Exp \left\{\ \times\ Exp\right\}_1 \\
\quad\quad --------- \\
\quad \mid \mathbb{P}\, Exp \\
\quad\quad --------- \\
\quad \mid Exp\ Infix\ Exp \qquad \mathrm{L} \\
\quad\quad --------- \\
\quad \mid Exp\ Exp \qquad \mathrm{L} \\
\quad\quad --------- \\
\quad \mid Prefix\ Exp \\
\quad\quad --------- \\
\quad \mid Exp\ Postfix \\
\quad\quad --------- \\
\quad \mid ClassHierarchy \\
\quad \mid Exp\ .\ Name \\
\quad \mid Name \left[GenActuals\right] \\
\quad \mid \{\ \left[ExpList\right]\ \} \\
\quad \mid \{SchemaText \left[\ \bullet\ Exp\right]\ \} \\
\quad \mid \varnothing \\
\quad \mid (Exp \left\{, Exp\right\}_1\ )
\end{array}
$$

$$\mid \langle \left[\mathit{ExpList} \right] \rangle$$

$$\mid \mathit{Number}$$

$$\mid \mathrm{self}$$

$$\mid (\mathit{Exp})$$

$$\mathit{LetDefs} ::= \mathit{LetDef} \left\{ ; \ \mathit{LetDef} \right\}$$

$$\mathit{LetDef} ::= \mathit{Name} \mathrel{\widehat{=}} \mathit{Exp}$$

$$\mathit{ClassHierarchy} ::= \left[\downarrow \right] \mathit{ClassDes}$$

$$\mathit{ClassDes} ::= \mathit{ClassName} \left[\mathit{GenActuals} \right] \left[\mathit{Renaming} \right]$$

$$\mathit{ClassName} ::= \mathit{Name}$$

$$\mathit{GenActuals} ::= [\mathit{ExpList}]$$

$$\mathit{Renaming} ::= [\mathit{RenameList}]$$

$$\mathit{Rename} ::= \mathit{Name} / \mathit{Name}$$

The syntax does not define the details of the productions for

$$\mathit{Infix}, \mathit{Prefix}, \mathit{Postfix}, \mathit{Number}, \mathit{Name}, \mathit{Sep}.$$

Elision of end-of-line semicolons, end-of-line conjunction symbols and ‘such that’ bars in 2-D structures without predicates, is not shown in the syntax.