

# Specifying Processes with Dynamic Life Cycles

Rick van Rein

University of Twente, Dept INF  
PO Box 217, NL-7500 AE Enschede, the Netherlands  
`vanrein@cs.utwente.nl`

**Abstract.** We propose an alternative notation and semantics for process models in object analysis, to resolve problems with current diagram languages. Our dynamic life cycles are communicating state diagrams. Our life cycles support polymorphic creation, and they are straightforwardly composed. We provide an operational semantics, and demonstrate how to interact with a system of life cycles.

**Keywords.** object orientation, polymorphic creation and deletion, conceptual modelling, communication, state diagrams, life cycles, process algebra, component composition.

## Introduction

Modern systems analysis and conceptual modelling is performed with object oriented methods, which offer rich notation to capture the constraints on implementations. An often-used notation is OMT, increasingly replaced with UML. The most-used models in these languages are *class diagrams* (which express classes with associations between them) and *state diagrams* (which express dynamic aspects of classes).

An example of a class diagram in OMT notation is found in figure 1, which shows a library (example taken from the analysis course on Catalysis [5]; Catalysis is an analysis and design method exploiting the notation of UML). This library registers *books*, which may be lent by *members*. A *book* has a *title*, against which *reservations* can be made.

Class diagrams have evolved from entity relationship diagrams [7], which explains why they represent data aspects but no process aspects. A result of the emphasis on data aspects in object oriented analysis is that associations lack a deeper meaning than “two connected classes.” Many designs do not even name associations!

In the design in figure 1, the *lentby* association optionally refers to a *member*, but no knowledge is added on when a *member* is found over that association. Furthermore, the meaning of the association can only be found by interpreting the word ‘lentby’ — a common source of misunderstanding. Similarly, the class diagram does not reveal whether the *lentby* and *heldby* links can lead to an object at the same time.

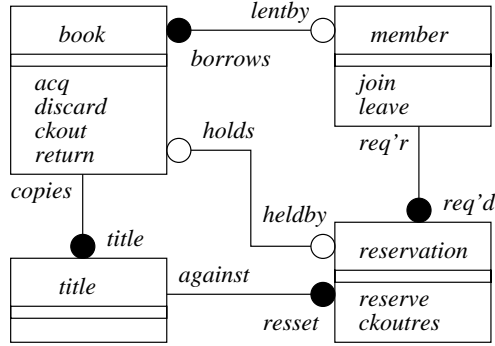


Fig. 1. Class diagram for a library system.




Fig. 2. State chart for a book in the library.

To express such dynamic aspects of the design in an object oriented way, a state diagram is usually drawn, like the one in figure 2, and some methods (like Catalysis) add mutually exclusive predicate expressions to the states, to hold for all *b:book*:

$$\begin{aligned}
 b.Available &= (b.lentby = nil) \wedge (b.heldby = nil) \\
 b.Borrowed &= (b.lentby \neq nil) \wedge (b.heldby = nil) \\
 b.Onhold &= (b.lentby = nil) \wedge (b.heldby \neq nil)
 \end{aligned}$$

This *does* answer our previous questions. A caveat here is the use of the value *nil*, which always introduces danger of loss of referential integrity; luckily, the relation between *nil* values and states makes it possible to check the models for *nil* dereferencing.

The state diagram in figure 2 demonstrates another place where a predicate expression can be of use, namely when the *return* operation is applied to a *book*; in this case, some condition (mentioned as *res* here) specifies whether or not the *book* is put on hold for some *reservation*. The disadvantage of this approach is that the *book* must be aware of the existence of *reservations*. It would be better for extensible designs when *books* are just entities with an *Available* and a *Borrowed* state, processing *ckout* and *return* operations, and not minding about *reservations*. Unfortunately this is not possible because a *ckout* for a reservation (distinguished here by naming it *ckoutres*) poses additional constraints and also terminates the resolved *reservation*.

Another class of process related problems in object oriented models is how to express synchronisation; for example, not enabling a *member* to *leave* the library (membership) before his *lends* set is empty. These problems are solved with preconditions on the *leave* operation in Catalysis.

Predicate logic expressions are powerful tools for object modelling, but not without a downside; free-form quantifications over sets make automatic model checking very hard and in general impossible since all possible runtime occurrences of sets cannot be foreseen at design time. Design specific ('free form') predicate expressions introduce the need for runtime checking, which has a negative impact on runtime performance.

The remainder of this paper is dedicated to introducing an (object) process model that contains more information than current life cycle models; the presentation is structured as follows: section 1 introduces and defines the concepts of life cycles by showing solutions to typical problems; section 2 shows a number of practical applications of the life cycle concept; finally, we draw conclusions and look forward to future research.

This work has been performed in the scope of the Quantum project, in which Compuware's UNIFACE lab and the University of Twente cooperate. UNIFACE is a leading component-based development tool for mission-critical applications.

*Our approach.* In this paper we introduce our notion of (*dynamic*) *life cycles*, as a variation on state diagrams. Life cycles can be used to express the process view on a class or component. They are capable of communicating with each other, and to create and destruct themselves on the fly. Life cycles pass around identities of life cycle instances (called *lives*, the process view on objects). We choose not to model nesting, direction, or fixate the number of participants for operations; instead, we model an event-like scheme with the possibility to extend the set of event participants in several ways. Our events interact synchronously, so that lives have the capability to block events.

We define the operational semantics for our life cycles in the functional programming language Miranda<sup>1</sup> [3]; it is a precise, high-level language that hardly forces us to over-specify, it is executable and relatively easy to read.

## 1 The Concept of Life Cycles

This section introduces the concept of (dynamic) life cycles, by showing which problems exist, and how they are solved by life cycles. Each problem is introduced with an example, and guides the definition of the graphical syntax and an operational semantics for life cycles.

### 1.1 Undirected Events

*Problem.* Events mentioned in state diagrams can be related to 'the real world' in a number of ways. State charts by Harel [9] and in UML [16] usually have a

<sup>1</sup> Miranda is a trademark of Research Software Ltd.

notion of events coming from outside the state charts' domain. To provide state charts with the ability of initiating further actions, they can send out events. This approach is also taken in ROOM [19] and in our own protocol checking work [17].

This approach is not free of problems. Firstly, it is often debatable [12] [2] what the semantics are if a multitude of events arrives at the same time, and whether a nesting of events is allowed. It seems that no reasonable choice covers all cases of practical interest. We believe that the problem is that the event-sends-event structure reveals too many details of implementation in programming languages, notably C++.

*Discussion.* We believe that this problem is caused by the imperative idea of events calling each other. To model all aspects of these calls well, the models either contain a lot of complicating detail, or fail to be sufficiently precise for verification by automatic tools. In contrast, process-algebras [1] are simpler and express similar information. In these models however, events are just atomic actions that communicate, with no distinction between caller and callee.

*Solution.* Our choice is to simplify the semantics of events in state charts, by removing knowledge of the direction of event communication. Instead, we observe something happening in the 'real world' (what we will henceforth call an *event occurrence*), and our life cycles define *events*, which are system-perspectives of event occurrences.

We define an *event* as a name plus a list of parameters (introduced in subsection 1.4). We define event occurrences *evocc* as a set<sup>2</sup> of events. This makes it possible to model complicated event occurrences that can be matched by different events in different places:

$$\begin{aligned} \text{event} &\equiv (\text{name}, [\text{param}]) \\ \text{evocc} &\equiv [\text{event}] \end{aligned}$$

Based on these definitions, we define a *lifecycle* as a triple containing a name, a set of states (which we define as simple names, unique within a *lifecycle*), and a set of *transitions* (which we introduce in subsection 1.6):

$$\text{lifecycle} \equiv (\text{name}, [\text{state}], [\text{transition}])$$

Where we leave the types *name* and *state* unspecified, and defer the definition of *transition* to later.

*Example.* In figure 3, we show a simple system comprising of two life cycles. An event occurrence in this system is  $[acq]$ , which matches the event *acq* in *book* but is ignored in *members*. Another event occurrence is  $[ckout]$ , which is matched by an event *ckout* in both *member* and *book* life cycles and which therefore is a

<sup>2</sup> Miranda does not support sets as language primitives; we will therefore use lists without taking order or multiplicity of elements into account.

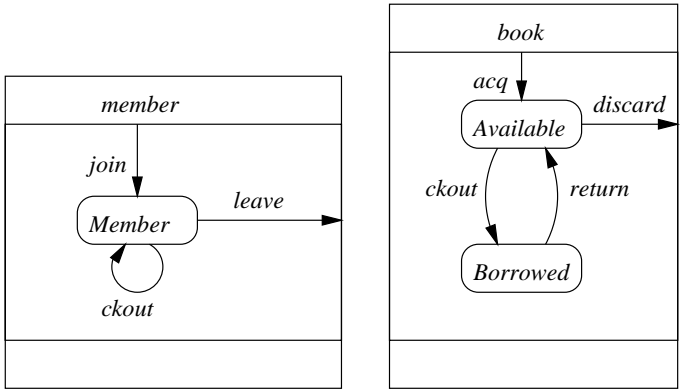


Fig. 3. Library in normal state diagram notation.

step that the two life cycles make at the same time. Though not useful in this application, it would be correct to consider  $[acq, leave]$  as an event occurrence, which is matched by  $acq$  in *book* and by  $leave$  in *member*. Instead of a nested calling structure between events, the events communicate by observing the same event occurrences at the same time.

### 1.2 Dynamic Communication Structures

*Problem.* Modern software development aims for modular (or ‘component-based’) development, in the hope that modules can be composed after they have been separately built. This makes it necessary to construct modules that can dynamically modify their patterns of interaction with the environment.

A problem encountered in implementing this wish is the strict pairing of caller to callee in most implementation languages. This style of specification makes it hard to add or remove a party without adding or removing its communication partners at the same time.

*Discussion.* An elegant and general scheme for communication is one where *any number* of parties may participate in an event. When a party wants to participate in an event, it cannot be ignored just because it currently is in a wrong state for participation. So, when a life cycle (like *book* in figure 3) contains the operation *return*, it can block that operation on other life cycles if it is not yet ready to participate in it. However, if a life cycle (like *member*) does not contain the *return* event, it simply ignores occurrences of that event.

The ideas of synchronisation (or ‘blocking’) described here are well known in concurrency theory. It is found in LOTOS [6] [4] as  $\parallel$ , and in ACP [1, §4.2], CSP [10] and CCS [14] as  $\parallel$ . Collaborative processing of events is known in programming languages as *rendez-vous synchronisation*. It is found in the languages Ada [11] and in Occam<sup>3</sup> [13]. Furthermore, there seems to be a correspondence between event occurrences and distributed transactions [21].

<sup>3</sup> Occam is a registered trademark of INMOS Limited.

*Solution.* To be able to talk about acceptance or blocking of event occurrences, we introduce a small algebra in which the interesting values are:

- *Ignore* to denote that an event (occurrence) is ignored;
- *Match*  $xs$  to denote that an event (occurrence) is accepted provided that the unifications in  $xs$  are established;
- *Block* to denote that an event (occurrence) is blocked and therefore may not take place.

The algebra provides a type *acceptance* with these values and a *compose* operation to compose lists of such values:

$$\begin{aligned}
 \text{acceptance} &::= \text{Ignore} \mid \text{Match} [(arg, arg)] \mid \text{Block} \\
 \text{isIgnore} &= (= \text{Ignore}) \\
 \text{isMatch}(\text{Match } b) &= \text{True} \\
 \text{isMatch } x &= \text{False} \\
 \text{isBlock} &= (= \text{Block}) \\
 \text{compose} &:: [\text{acceptance}] \rightarrow \text{acceptance} \\
 \text{compose } xs &= \text{Block}, & \text{if } \exists x \in xs @ \text{isBlock } x \\
 &= \text{Match } [v \mid (\text{Match } vs) \leftarrow xs; v \leftarrow vs], & \text{if } \exists x \in xs @ \text{isMatch } x \\
 &= \text{Ignore}, & \text{otherwise}
 \end{aligned}$$

Using this *acceptance* algebra, we can define a function that indicates how two events match. This function uses the notion of parameters in events, which are introduced in subsection 1.4; in short, a parameter can either be an *In* or *Out* value with the usual meaning, or it can be a *Force* value which means that an event with a different value for that parameter *Blocks* if it matches on the non-*Forced* parameters.

The *eventmatch* function returns *Ignore* if there is a discrepancy between the types of the two events. Note that *Out* parameters are treated as part of the event name. If the events do not *Ignore* each other, they should *Match* or *Block*; the latter case occurs when a *Force* parameter (introduced in subsection 1.4) has a wrong value. When *Match* is returned, variable bindings due to parameters are returned as well. So:

$$\begin{aligned}
 \text{eventmatch} &:: \text{event} \rightarrow \text{event} \rightarrow \text{acceptance} \\
 \text{eventmatch } (n, p) (n', p') &= \text{Ignore}, & \text{if } n \neq n' \vee \#p \neq \#p' \\
 &= \text{Ignore}, & \text{if } \exists x \in \text{matchers} @ \text{mismatch } x \\
 &= \text{Block}, & \text{if } \exists x \in \text{unification} @ \text{blocking } x \\
 &= \text{Match } \text{unification}, & \text{otherwise}
 \end{aligned}$$

where

$$\begin{aligned}
 \text{matchers} &= \text{zip}(p, p') \\
 \text{mismatch}((\text{Out}, \text{Bound } v), (\text{Out}, \text{Bound } v')) &= v \neq v' \\
 \text{mismatch}((d, w), (d', w')) &= \text{False} \\
 \text{unification} &= [(w, w') \mid ((d, w), (d', w')) \leftarrow \text{matchers}] \\
 \text{blocking}((\text{Bound } v), (\text{Bound } v')) &= v \neq v' \\
 \text{blocking}(w, w') &= \text{False}
 \end{aligned}$$

*Example.* With the above definition of *eventmatch*, it is possible to verify the examples in the previous subsection. Since we have not yet considered parameters, we will set the lists of parameters to  $\square$ , the empty list. The event *acq* is modelled formally as  $(\text{“acq”}, \square)$  and we find that  $\text{eventmatch}(\text{“acq”}, \square)(\text{“acq”}, \square) = \text{Match } \square$  and  $\text{eventmatch}(\text{“acq”}, \square)(\text{“ckout”}, \square) = \text{Ignore}$ . We generalise this notion to the matching with event occurrences in subsection 1.6.

The problem of allowing dynamically changing numbers of collaborators on an event occurrence can be solved with the *eventmatch* definition: Those events that *Ignore* an occurring event will not take part in it; those that *Block* cause the whole event occurrence to *Block* (thanks to the *compose* operator) and if this does not happen, all events that *Match* take part in the occurring events.

### 1.3 Creation and Destruction of Lives

*Problem.* A running life cycle system must support instances, which we call *lives*: some *books* in our library will be *Borrowed*, while others are *Available*. To keep track of different books, life cycles must be instantiated to *lives*. Only for a life is it sensible to speak of its current state, and a life is what responds to event occurrences, or blocks them.

Some common problems related to creation and destruction of lives or objects are polymorphic creation of a new instance (often solved with factories [8]) and when to destruct a life or object (often solved with garbage collection [22]).

*Discussion.* We prefer to solve creation and deletion of lives without explicit operations. We prefer creation and destruction to be a response of a life cycle system to a plain event occurrence, because this hardly requires additional life cycle concepts.

Process algebras such as ACP [1, table 35], which model concurrent processes, silently remove a finished process; this is part of the design of the concurrency operators. There is also a process algebra, namely the  $\pi$ -calculus [15], which introduces a similarly elegant creation construct: a process  $P$  may be prefixed with a replication operator, with the interesting property that  $!P = P||!P$ , which is any number of  $P$  processes interleaved. This appears like an infinite-sized pool of instances of  $P$  being ready to get created at the first communication with one of the  $P$ s.

The  $\pi$ -calculus approach to process creation is generally considered elegant because it eliminates the requirement of explicit creating operations. Creation is just an effect of normal communication in the  $\pi$ -calculus.

*Solution.* To handle instances of life cycles, which we call *lives*, we define a *life* to be a quadruple with its *lifecycle*, an identity that is unique for that life cycle, a current *state* taken from the life cycle, and a binding relation for currently known variables (treated in subsection 1.4):

$$\text{life} \equiv (\text{lifecycle}, \text{lifeID}, \text{state}, \text{binding})$$

The type *lifeID* is only required to support an equivalence relation; in our examples, we assume that *lifeID* is represented as *num*. This paper works only with

life cycles with one creating transition, and these are mutually exclusive, so that for each life cycle, at most one instance is created during any event occurrence.

We define a global system state *sys* as a set of life cycles and a set of lifes that instantiate those life cycles:

$$sys \equiv ([lifecycle], [life])$$

We define creation and deletion as actions on such a *sys*. Later restrictions will ensure that only one instance at a time can be created for each life cycle. We already exploit this knowledge here and create one *life* for every *lifecycle* in a system, exploiting a special state *nirwana* that we reserve to represent the state of a life before its creating transition(s) or after its destroying transition(s):

$$\begin{aligned} nirwana &:: state \\ nirwana &= \text{“Nirwana”} \\ inNirwana &:: life \rightarrow bool \\ inNirwana(lc, lid, st, b) &= (st = nirwana) \end{aligned}$$

The *nirwana* state is useful, because it enables the creation of non-existent lifes in a pre-natal state in the following *birth* function, and destroys outlived and never-used lifes in the *death* function:

$$\begin{aligned} birth, death &:: sys \rightarrow sys \\ birth(lcs, ls) &= (lcs, ls + [(lc, nid\ lc, nirwana, [(this, nid\ lc))] \mid lc \leftarrow lcs]) \\ &\text{where} \\ &\quad nid\ lc = id\_not\_in\ [lid \mid (lc', lid, y, z) \leftarrow ls; lc' = lc] \\ death(lcs, ls) &= (lcs, ls - [l \mid l \leftarrow ls; inNirwana\ l]) \end{aligned}$$

We do not specify *id\_not\_in*, except that for all *xs* it holds that  $(id\_not\_in\ xs) \notin xs$ .

*Example.* The *birth* and *death* functions are applied just before and after a basic step of a life cycle system. This means that for every lifecycle, there is always a life ready to be instantiated in response to an event occurrence. Similarly, a life can engage in an event occurrence which terminates its life, and it would be implicitly cleaned up by the *death* function.

To demonstrate the conceptual simplicity of our approach to creation, imagine adding a (redundant) constraint to our library, to express that every *ckout* of a *book* must be followed by a *return* of that *book*. This can be done by adding the life cycle in figure 4. This creates a new life on the fly, as a side-effect of the occurrence of a plain *ckout* event. The *ckout* event receives no special treatment anywhere else than on the creating arrow in this new life cycle; deletion is similar.

## 1.4 Parameters and Variables

*Problem.* If we instantiate numerous lifes from the *book* life cycle, we do not wish them to jointly transit from the *Available* state to the *Borrowed* state, but






Fig. 4. The checkout life cycle.

rather one at a time. However, it is at the same time useful to let events be globally visible, to enable any life to participate in any event occurrence.

*Discussion.* We need to give every life its own identity, with a *this* notation, just like the keyword representing the ‘current object’ in object-oriented programming languages. The *this* keyword makes it possible to distinguish lives from each other; but the events should also distinguish between different lives involved, rather than life cycles. Therefore, we introduce parameters on events.

Normal object-oriented programming languages often support *overloading* of operations, where two operation names are only equivalent when the number and types of parameters somehow match. We take this same approach, and go one step further, by not only looking at the types of parameters, but also at the identities of the parameters. This is like multi-methods, defined from a process perspective.

To make it possible to give out parameters other than *this*, a life cycle can define variables to hold references to (other) lives. These variables can be read or written as event parameters.

*Solution.* Every life can define a set of variables, which hold references to (other) lives. These values are passed around in parameters to events and event occurrences. Two important properties of parameters are their direction and whether or not they are (already) bound to a value.

The direction of parameters may be:

- *In*, denoted with *?* graphically, for variables whose value is read from the parameter. The value for these parameters is not constrained by this parameter.
- *Out*, the graphical default, for variables that are exposed to a match along with the name of the event. The third clause of *eventmatch* exists because of this matching on *Out* parameters.
- *Force*, denoted with *!* graphically, to enforce that an event’s parameter matches a variable’s value. An example of this will be introduced in subsection 1.5. If the rest of an event *e* matches but this parameter does not, then *e* is *Blocked*.

The parameter may either be an *Unbound* reference to a variable (mentioning that variable by name) or it may be *Bound* to some *lifeID* value:

$$\begin{aligned} param &\equiv (dir, arg) \\ dir &::= In \mid Out \mid Force \\ arg &::= Unbound \ var \mid Bound \ lifeID \\ var &\equiv name \end{aligned}$$

Variable names are unique within a life cycle. Furthermore, every life cycle defines a variable named “**this**”, that already holds the *lifeID* of a life on creating transitions; all other variables can only get values by obtaining them as *In* parameters from some event occurrence:

$$\begin{aligned} this &:: var \\ this &= \text{“this”} \end{aligned}$$

A life cycle definition only contains *Unbound* parameters, but after a life is instantiated, variables are filled in. We define a *binding* to be a mapping from *Unbound* to *Bound* parameters; this is the way variable values are represented in a *life*. The *bind* function applies a *binding* by mapping it over a *transition*, whose introduction we defer to subsection 1.6:

$$\begin{aligned} binding &\equiv [(var, lifeID)] \\ bind &:: binding \rightarrow transition \rightarrow transition \\ bind \ bnd \ (f, t, eves) &= (f, t, map \ (eventbind \ bnd) \ eves) \\ \text{where} \\ eventbind \ bnd \ (nm, parms) &= (nm, map \ (parmbind \ bnd) \ parms) \\ parmbind \ ((nm, val) : nms) \ (d, Unbound \ nm) &= (d, Bound \ val) \\ parmbind \ (nv : nms) \ x &= parmbind \ nms \ x \\ parmbind \ y \ x &= x \end{aligned}$$

A variable *v* may only be used as an *Out* or *Force* parameter in transitions leaving a state *s* if either *v* is *this* or *v* is mentioned at least once as an *In* parameter on each possible path leading to state *s*. This is a syntactical constraint on life cycles.

*Example.* Figure 5 presents a redesign of the library of figure 3, where the events have parameters. For example, the event occurrence *ckout(b,m)* communicates with *book b* and *member m*; therefore, a single *book b* is checked out, and other books remain in the state in which they were; also, if some *book b'* is in *Borrowed* state, then it does not block the *ckout(b,m)* event occurrence, provided that *b'* differs from *b*.

## 1.5 Enforcing Synchronisation

*Problem.* Imagine extending the library in figure 5 with reservations (on *books* rather than *titles* for simplicity). Then, whenever a reserved *book* is returned,

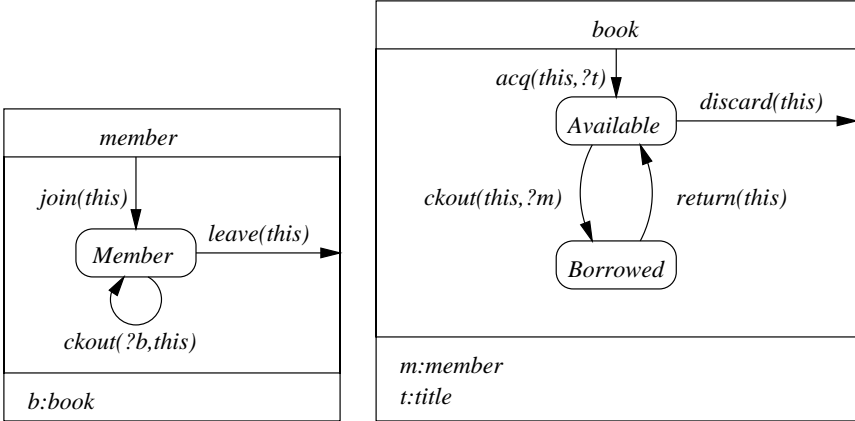


Fig. 5. Two library life cycles with parameters.

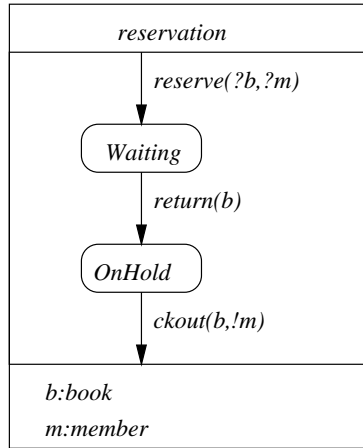
it must be put on hold. In many object methods, this is done by splitting the *return* transition like in figure 2. In our life cycle approach, we prefer to exploit communication to add a constraint, stating that the next *member* to check out the *book* *b* must be the *member* for whom *b* was put on hold.

*Discussion.* With only *In* and *Out* parameters, life cycles could not express this kind of ‘identity enforcement.’ An *In* parameter matches with anything, and an *Out* parameter simply *Ignores* event occurrences that have different identities for that variable. Therefore, we need to introduce another kind of parameter, which disallows (or *Blocks*) event occurrences with non-matching parameters.

*Solution.* We will graphically annotate such parameters with a *!* which is pronounced ‘(en)force.’ In our formal models, we use the direction *Force* for these parameters. If an event occurrence *eo* matches an event on name, parameter count and all *Out* parameters, then there still is the question whether *eo* should be *Matched* or *Blocked*; the latter is chosen in case a *Force* parameter has different value from the value in *eo*. This is what has been formalised in clause 5 of the definition of *eventmatch*.

*Example.* Figure 6 defines the *reservation* life cycle. This life cycle is created when a *reserve(?b,?m)* event occurs, and awaits the *return* of *book* *b*. When this book is returned, the *reservation* life claims it, and adds a constraint on the next *ckout* of *b*, stating that it may *only* be checked out by *member* *m* who made the *reservation* initially. The result is a dynamic extension to the communication structure which deals with the added constraint due to the *reservation*; the *book* need not be aware of this additional constraint.

Note how the asymmetry between the book and the member in *ckout* corresponds to a conceptual asymmetry: *member* *m* should be able to check out other *books* than *b*, but *book* *b* may not be checked out by another *member* than *m*.



**Fig. 6.** Enforcing a certain *member* parameter value.

## 1.6 Multiple Perspectives on Event Occurrences

*Problem.* In the previous solution, we introduced a *reservation* on *books*. It is better to reserve *titles* than *books*, to support the situation where two copies (*books*) of the same *title* are owned by the library. Such a *reservation* life cycle that awaits a *title* cannot be made without requiring the title in the parameter list of the events that make a book available, thus also on *return*. The *title* has thus far not been a parameter to the *return* event, and it would be a demonstration of bad extensibility of life cycles if we would add it now.

*Discussion.* This is a common problem; from one perspective, there should be certain parameters to an event occurrence, and from another perspective there should be other parameters. This is caused by different perspectives on the same event occurrence. From the perspective of a *book*, only the *book* returned is of interest. From the perspective of a *reservation*, the *title* is also of interest.

We want to support observing event occurrences from different perspectives, much like in step failures semantics [20]. We therefore allow *multiple events per transition*. Each of these events can represent a different perspective on the transition. To ensure that all perspectives occur at the same time, the events on a transition must all occur together.

Note how this approach replaces that of nested states, as used in state charts. Nested states have a hierarchical ordering, and that hierarchy is for one perspective only. Another perspective usually shows up as a cross-level transition, which indeed is part of the syntax of state charts. We therefore believe that an extension of life cycles with hierarchical states would *weaken* the concept of life cycles rather than improve it.

*Solution.* A transition in a life cycle is annotated with any number of events:

$$\begin{aligned} \text{transition} &\equiv (\text{state}, \text{state}, [\text{event}]) \\ \text{transition} &:: \text{lifecycle} \rightarrow \text{state} \rightarrow [\text{transition}] \\ \text{transition} &(\text{n}, \text{x}, \text{ts}) \text{ from} = [(\text{f}, \text{t}, \text{es}) \mid (\text{f}, \text{t}, \text{es}) \leftarrow \text{ts}; \text{f} = \text{from}] \end{aligned}$$

Now that we have seen that event occurrences as well as transitions are represented with sets of events, it is possible to define when they match; informally, this is true when for every event in the transition, there is a matching event in the event occurrence:

$$\begin{aligned} \text{evocmatch} &:: \text{evocc} \rightarrow \text{transition} \rightarrow \text{acceptance} \\ \text{evocmatch } \text{eo} (\text{f}, \text{t}, \text{es}) &= \text{Block}, && \text{if blocked} \\ &= \text{compose} [\text{eventmatch } \text{e}' \text{ e} \mid \text{e}' \leftarrow \text{eo}; \text{e} \leftarrow \text{es}], \\ &\text{where} && \text{otherwise} \\ \text{matchdom} &= [e \mid e \leftarrow \text{es}; \exists x \in \text{eo} @ (\text{isMatch} \circ \text{eventmatch } e) x] \\ \text{blocked} &= [] \neq \text{matchdom} \neq \text{es} \end{aligned}$$

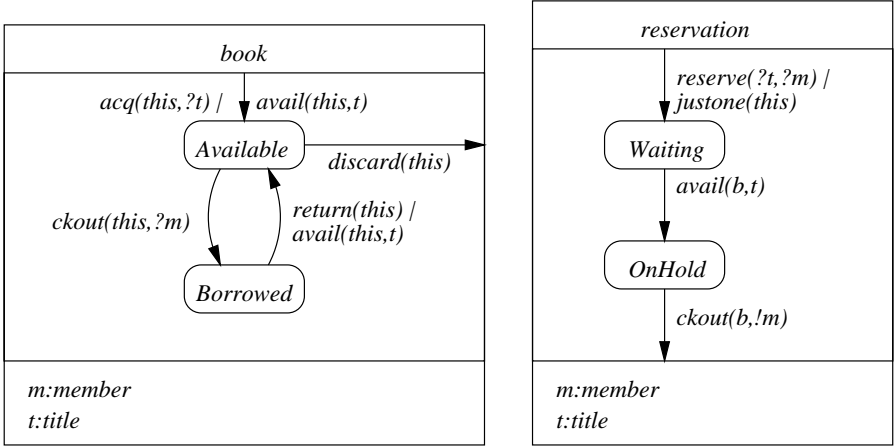
In general, a *match* is found if no inconsistencies exist in the required *Match* unifications:

$$\begin{aligned} \text{match} &:: \text{evocc} \rightarrow \text{transition} \rightarrow \text{acceptance} \\ \text{match } \text{eo } \text{tr} &= \text{consistent}(\text{evocmatch } \text{eo } \text{tr}) \\ &\text{where} \\ \text{consistent } \text{Ignore} &= \text{Ignore} \\ \text{consistent}(\text{Match } \text{xs}) &= \text{Match } \text{xs}, \\ &\text{if and } [\text{a}=\text{a}' \mid (\text{x}, \text{a}) \leftarrow \text{xs}; (\text{x}', \text{a}') \leftarrow \text{xs}; \text{x} = \text{x}'] \\ \text{consistent } \text{x} &= \text{Block} \end{aligned}$$

*Example.* To notate multiple events on one transition, we intersect the transitions with a comma, pronounced ‘and.’ For example, in the life cycle for a *book*, it is sensible to introduce an event *avail(this,t)* along the transition that now holds the event *return(this)* only. The *avail* event can be caught in the *reservation* life cycle. Moreover, from the perspective of books becoming available, we can see the acquisition of a new *book* as an *avail(this,t)* event too. The resulting life cycles for *book* and *reservation* are therefore as given in figure 7. It is sufficient to *mention* the existence of a *title*, we need not define its life cycle.

## 1.7 Step Semantics

This subsection is dedicated to describing precisely what the semantics of a step of a complete life cycle system are. The general approach is to first find out what transitions a single life can make given some event occurrence, then what happens if these are applied to a single life, and then map this approach over all life cycles in the system.



**Fig. 7.** The use of multiple events per transition.

The *lifesteps* function lists every possible *transition* that a life can engage into, given its current state and an event occurrence. The result lists the *acceptance* for every possible *transition*. The result is a set (implemented here as a list); an empty set means that no non-blocking transitions are possible; one result means that a unique transition can be made; multiple results means that a non-deterministic choice can be made between the outcomes.

$$\text{lifesteps} :: \text{evocc} \rightarrow \text{life} \rightarrow [(\text{transition}, \text{acceptance})]$$

$$\text{lifesteps eo } (lc, lid, st, b) = \text{matched} \ \underline{\text{else}} \ \text{ignored}$$

where

$$\text{transitions} = [(\text{tr}, \text{match eo tr}) \mid \text{tr} \leftarrow \text{map } (\text{bind } b) (\text{transfrom } lc \ st)]$$

$$\text{nonblocked} = [(\text{tr}, \text{acc}) \mid (\text{tr}, \text{acc}) \leftarrow \text{transitions}; \neg \text{isBlock acc}]$$

$$\text{matched} = [(\text{tr}, \text{acc}) \mid (\text{tr}, \text{acc}) \leftarrow \text{nonblocked}; \text{isMatch acc}]$$

$$\text{ignored} = [(\text{tr}, \text{acc}) \mid (\text{tr}, \text{acc}) \leftarrow \text{nonblocked}; \text{isIgnore acc}]$$

$$\square \ \underline{\text{else}} \ \text{ys} = \text{ys}$$

$$\text{xs} \ \underline{\text{else}} \ \text{ys} = \text{xs}$$

This makes it possible to construct a function *steplife* that constructs a list of possible future *lives* from a current *life* and an *evocc* happening to it:

$$\text{steplife} :: \text{evocc} \rightarrow \text{life} \rightarrow [\text{life}]$$

$$\text{steplife eo } l = \text{map } (\text{actualstep } l) (\text{lifesteps eo } l)$$

With the *actualstep* function, a transition is applied to a life; this influences the state in which the life resides. Furthermore, old bindings *bnd\_old* are removed as far as they are read by this step, and new bindings *bnd\_new* are added instead for read-in variables:

$$\text{actualstep} :: \text{life} \rightarrow (\text{transition}, \text{acceptance}) \rightarrow \text{life}$$

$$\text{actualstep } (lc, lid, st, bnd) (\text{tr}, \text{Ignore})$$

$$\begin{aligned}
&= (lc, lid, st, bnd) \\
&actualstep (lc, lid, st, bnd) ((f, t, es), Match ms) \\
&= (lc, lid, t, mkset((bnd \text{ --- } bnd\_old) \text{ ++ } bnd\_new)) \\
&\text{where} \\
&bnd\_old = [(var, val) \mid (var, val) \leftarrow bnd; \\
&\quad \exists x \in ms @ ((Unbound var =) \circ snd) x] \\
&bnd\_new = [(var, val) \mid (Bound val, Unbound var) \leftarrow ms]
\end{aligned}$$

A life can block the execution of an *evocc* that it can accept in some possible future; this is a consequence of the dynamic communication principle explained in subsection 1.2. The function *futureblock* determines if the given *sys* blocks the given *evocc* because of some future transition of which all variables are currently bound:

$$\begin{aligned}
&futureblock :: evocc \rightarrow sys \rightarrow bool \\
&futureblock eo (lcs, ls) = \exists x \in ls @ futureblockinglife x \\
&\text{where} \\
&futureblockinglife l = \exists x \in futurematches l @ ((\neg) \circ isIgnore) x \\
&futurematches l = [acc \mid (tr, acc) \leftarrow futurelifesteps [nirwana] eo l; \\
&\quad allbound tr] \\
&allbound(f, t, es) = [var \mid (nm, parms) \leftarrow es; \\
&\quad (d, Unbound var) \leftarrow parms; d \neq In] = []
\end{aligned}$$

This uses a function *futurelifesteps*, which lists, once again for each possible *transition*, the possible *acceptance* value, only this time taking the future after that step into account in the *acceptance* value. This blocks on events that are *Ignored* or *Matched* in the current state, but that *Block* in some future state.

$$\begin{aligned}
&futurelifesteps :: [state] \rightarrow evocc \rightarrow life \rightarrow [(transition, acceptance)] \\
&futurelifesteps states eo l = [(tr, compose(acc(tr, a) : accs(tr, a))) \mid \\
&\quad (tr, a) \leftarrow stepstodo] \\
&\text{where} \\
&stepstodo = lifesteps eo l \\
&stepsdone((f, t, es), a) = \\
&\quad filter otherstate [actualstep l ((f, t, es), Match [])] \\
&otherstate(lc, lid, s, b) = s \notin states \\
&acc = compose \circ map snd \circ concat \circ map(futurelifesteps (states' l) eo) \\
&\quad \circ stepsdone \\
&accs(tr, a) = [], \quad \text{if } \# \text{ states} = 1 \\
&\quad = [a], \quad \text{otherwise} \\
&states'(lc, lid, s, b) = s : states
\end{aligned}$$

When a *sys* makes a step, every *life* decides independently which of its possible futures to select; therefore, the whole *sys* has as its possible futures the cartesian product of the possible futures of all its *lives*. Note how we apply the functions

*birth* and *death* to obtain the desired effect of creation and destruction as side-effects of communication:

$$\begin{aligned}
 \text{step} &:: \text{evocc} \rightarrow \text{sys} \rightarrow [\text{sys}] \\
 \text{step } eo \text{ sys} &= [], && \text{if } \text{futureblock } eo \text{ (birth sys)} \\
 &= ((\text{map } \text{death}) \circ \text{sysstep} \circ \text{birth}) \text{ sys}, && \text{otherwise}
 \end{aligned}$$

where

$$\begin{aligned}
 \text{sysstep}(lcs, ls) &= [(lcs, ls') \mid \\
 &\quad ls' \leftarrow \text{cartesian}(\text{map } (\text{mkset} \circ \text{steplife } eo) \text{ ls})]
 \end{aligned}$$

The *step* function is the interface through which users of a life cycle system can make the system perform steps.

## 2 Interacting with a Life Cycle System

The previous section introduced our life cycle concept. In this section we demonstrate how commonly requested information is available in a life cycle system. We describe the demonstrations in this example in our semantics language Miranda, to achieve sufficient precision, also outside the boundaries of our life cycles.

### 2.1 A Component with a Book LifeCycle

The *book* life cycle in figure 7 is translated into the formal notation by introducing a number of transitions with sets of events on them, and building the *book* life cycle from it:

$$\begin{aligned}
 \text{book} &:: \text{lifecycle} \\
 \text{book} &= (\text{"book"}, [\text{nirwana}, \text{"Available"}, \text{"Borrowed"}], \\
 &\quad [\text{bookT1}, \text{bookT2}, \text{bookT3}, \text{bookT4}])
 \end{aligned}$$

where

$$\begin{aligned}
 \text{bookT1} &= (\text{nirwana}, \text{"Available"}, [(\text{"acq"}, [(Out, Unbound this), \\
 &\quad (In, Unbound \text{"t"})]), (\text{"avail"}, [(Out, Unbound this), \\
 &\quad (Out, Unbound \text{"t"})])]) \\
 \text{bookT2} &= \dots
 \end{aligned}$$

It is possible to apply event occurrences, including things like members joining the library, to this system (or component) of one life cycle; it simply ignores such event occurrences because no (future) transition of the *book* communicates with a *join* event. This is useful, since it means that an additional life cycle (or component) that does take the *join* event occurrence into account can process it. The *book* does not care about *joins*, but is not blocking them either; this is the idea of compositionality behind life cycles.

Assuming a library with a *book* life cycle present (and no *reservation* life cycle), it is possible to acquire a new book with event occurrences like *acq*:

$$\text{acq}, \text{ckout}, \text{return} :: \text{lifeID} \rightarrow \text{lifeID} \rightarrow \text{sys} \rightarrow [\text{sys}]$$



```

discard :: lifeID → sys → [sys]
acq b t = step [(“acq”, [(Out, Bound b), (Out, Bound t)]),
                (“avail”, [(Out, Bound b), (Out, Bound t)])]
ckout b m = ...
return b t = ...
discard b = ...

```

The *acq* event occurrence matches the *bookT2* transition in the *book* life cycle when the right values for *b* and *m* are filled in. Assume a library system with no instances: (*[book]*, []). Applying *acq* 1 3 to this system returns a list with one future library system, in which a book with *lifeID* 1 and title 3 is registered in state *In*. Applying *ckout* 1 7 to this result leads to a library system with that same book in *Borrowed* state.

```

exlib0 = ([book], [])
[exlib1] = acq 1 3 exlib0
[exlib2] = ckout 1 7 exlib1
nolibs = discard 1 exlib2
[exlib3] = return 1 3 exlib2
[exlib4] = discard 1 exlib3

```

It is impossible to *discard* that book from the resulting library *exlib*<sub>2</sub>; it should first be returned. This is the kind of synchronisation constraints that we strive for with life cycles, and which is inspired on the communication primitives in process algebras.

Note how the *step* function for the *ckout* and *acq* functions is used in the same way; the fact that *acq* constructs a new object is not visible here, meaning that we succeeded in hiding the creation (and destruction) of objects entirely.

## 2.2 Deriving Association Information

In the introduction of this paper, we argued that associations contain too little temporal information, and often require further constraints. In this subsection, we demonstrate how association information can be obtained from a life cycle system.

The vagueness of associations means that several definitions are thinkable; for example, in which states of the ‘from’ object (or life) the association is defined. We introduce the following helper function to define such associations:

```

assoc :: lifecycle → [state] → var → sys → [(lifeID, lifeID)]
assoc lc ss v (lcs, ls) = [(f, t) | (lc', f, s, b) ← ls; lc = lc';
                               s ∈ ss; (v', t) ← b; v = v']

```

Given a life cycle  $lc$  and the states  $ss$  in which the association should hold, and a variable  $v$  in  $lc$  (which must be defined in all the states  $ss$ ) points to the ‘to’ object (or life), this function returns a list of links (association instances) between ‘from’ and ‘to.’

An example application of this function is the *lentby* association from the introduction:

```
lentby :: sys → [(lifeID, lifeID)]
lentby = assoc book ["Borrowed"] "m"
```

For example,  $lentby\ exhib_1 = lentby\ exhib_3 = []$  but  $lentby\ exhib_2 = [(1, 7)]$ , denoting that in  $exhib_2$  a book 1 is *Borrowed* by member 7.

Note how life cycle models need no *nil* values to model absence of links; instead, process state information is used to say which variables are defined.

### 2.3 Global Process Analysis

Unlike UML state charts, our life cycles communicate in a precisely defined way. The semantics of life cycles thus expresses sufficient information to infer global system behaviour. Consequently, we see possibilities to perform a number of interesting global process checks on a system of life cycles. We intend to perform research on these process checks [18], in spite of the unbounded number of lifes that can be created.

*Deadlock avoidance.* One way to avoid deadlock is to ensure the following: *Given that the life cycle system  $s \in S$ , there must be a trace of event occurrences that leads from  $s$  to a system without lifes.* The smallest set  $S$  to avoid deadlock contains precisely those states that are reachable from a life cycle system without lifes. If a deadlock avoiding checker accepts a life cycle system, an important class of problems has been avoided; furthermore, in the attempts to come to such a system design, the checker can provide valuable feedback to the designer.

*Referential integrity.* In systems that deal with pointers (like variables in our life cycles), it is important to ensure referential integrity. In systems that use *nil* pointer values, it is desirable to avoid the possibility of runtime exceptions caused by *nil* dereferences. In general, these problems are hard to solve before runtime. For life cycles, this is simpler, since they specify global process knowledge more directly.

We intend to make a checker that ensures the following: *If a life  $l$  is in Nirwana state, it is not referenced anywhere.* In other words, in Nirwana state it is safe to delete a life. This resolves the debate [22] on whether to garbage collect or to explicitly delete, since the moment of deletion is precisely known.

*Non-determinism.* We explicitly allow non-determinism caused by more than one transition from the same life cycle state matching the same event occurrence. This is because non-determinism has a useful interpretation: something

outside the current scope of reference (e.g. some data-related issue) decides on the transition to follow. Checker tools must treat non-determinism as a choice that they cannot influence or count on. In our previous work on protocol checking [17], we used non-determinism in a similar way, namely to model choices made during the execution of an atomic action.

### 3 Conclusions and Future Work

This paper introduced a variation on state diagrams called (dynamic) life cycles. We presented a graphical notation and a precise semantics.

Our life cycles present some features that we did not find in comparable work: life cycles support polymorphic creation and deletion of new instances as a side-effect of normal steps. Life cycles systems are straightforwardly composed to larger systems. Life cycles can express dynamically changing synchronisation constraints. Life cycles allow redundancy and non-determinism in models.

Life cycles can express multiple perspectives on events as a means of dynamic binding; we expect this to be a good replacement for hierarchical state charts, whose precise meaning is not trivial.

Life cycles contain more information than traditional models. Associations can be inferred from life cycles. Information on which steps can be performed next, and which instances can act as parameters to those steps, can be inferred from life cycles. Finally, life cycles describe how the process aspects of a total system behave, which makes us believe that several global correctness properties become testable; think of deadlock avoidance and referential integrity.

In our future work, we hope to perform such global correctness checks in spite of the unbounded number of lives in a life cycle system. We intend to gain better understanding on how processes in workflow and implementation objects cooperate. We also wish to investigate the use of refinement in workflow models. We are also interested in relating this work to models like CSP and LOTOS to see if life cycles introduce new formal aspects (possibly, dynamically evolving communication alphabets). Our eventual goal with this work is to come to a simple yet powerful process notion and accompanying verification tools for object oriented models.

*Acknowledgements.* I wish to thank Maarten Fokkinga for providing constant feedback and input on the ideas behind life cycles. I wish to thank Compuware, in particular Wim Bast, Tom Brus and Edwin Hautus, for providing the challenging practical environment in which this work evolved.

### References

1. J.C.M. Baeten and W.P. Weijland. *Process Algebra*. Cambridge University Press, 1990.
2. M. von der Beeck. A comparison of Statecharts variants. In H. Langmaack, W.P. de Roever, and J. Vytöpil, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems*, pages 128–148. Springer, 1994. Lecture Notes in Computer Science 863.

3. R. Bird and P. Wadler. *Introduction to functional programming*. Prentice Hall, 1988.
4. T. Bolognesi and E. Brinksma. Introduction to the ISO specification language LOTOS. *Computer Networks and ISDN Systems*, 14(1):25–29, 1988.
5. D. D'Souza and A. Wills. *Catalysis: Practical Rigor and Refinement*. Addison-Wesley, 1998.
6. P.H.J. van Eijk, C.A. Vissers, and M. Diaz, editors. *The formal description technique LOTOS*. Elsevier Science Publishers B.V., 1989.
7. S.B. Elmasri, R. Navathe. *Fundamentals of database systems*. Benjamin/Cummings, 1994.
8. E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-oriented Software*. Addison Wesley, Reading, 1996.
9. D. Harel. Statecharts: A visual formalism for complex systems. *Science of Computer Programming*, 8:231–274, 1987.
10. C.A.R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
11. ISO, editor. *Ada 95 Reference Manual, Language and Standard Libraries*. 1994. ISO/IEC 8652:1995(E).
12. D. Latella, I. Majzik, and M. Massink. Towards a formal operational semantics of UML statechart diagrams. *Formal Methods for Open Object-Based Distributed Systems*, pages 331–347, 1999.
13. Inmos Ltd. *Occam 2 Reference Manual*. Prentice-Hall, 1988.
14. R. Milner. *A calculus of communicating systems*. LNCS. Springer-Verlag, 1980.
15. R. Milner. The polyadic  $\pi$ -calculus: a tutorial. Springer-Verlag, Aug 1991.
16. Rational Software Corporation. *UML Semantics*. Rational Software Corporation, 1997.
17. R. van Rein and M. Fokkinga. Protocol assuring universal language. *Formal Methods for Open Object-Based Distributed Systems*, pages 241–258, 1999.
18. R. van Rein and M.M. Fokkinga. Static checking of dynamic protocols. 1999. submitted to CONCUR'99.
19. B. Selic, G. Geullekson, and P.T. Ward. *Real-time Object-Oriented Modeling*. John Wiley & Sons, Inc., 1994.
20. D. Taubner and W. Vogler. Step failures semantics and a complete proof system. *Acta Informatica*, 27:125–156, 1989.
21. X/Open, editor. *Distributed Transaction Processing: Reference Model, Version 3*. Feb 1996.
22. B. Zorn. The measured cost of conservative garbage collection. *Software, Practice and Experience*, 23(7):733–756, July 1993.