

# A Comparison of LOTOS and Z for Specifying Distributed Systems \*

Iain Fogg  
Brian Hicks  
Andrew Lister  
Tim Mansfield  
Kerry Raymond  
Key Centre for Software Technology  
Dept. of Computer Science  
University of Queensland  
St Lucia QLD 4067  
Australia

10 May 1989

## Abstract

Although communications underpin the development of distributed systems, the same formal description techniques may not be equally appropriate for both problem domains. We compare two formal description techniques, the standardized LOTOS and the emerging Z, to evaluate their merits for the specification of distributed systems and distributed algorithms. Some of the aspects we examine include the underlying model of both FDTs, their support for the fundamental concepts of *process* and *communication*, their expressiveness in specifying the invocation source and sequence of operations, and their ability to specify liveness properties and other temporal constraints. Software tools are briefly investigated, but no attempt is made to examine available tools in-depth. We also describe our comparative case study; the specification of a distributed termination algorithm in LOTOS and Z. The reader is assumed to have some familiarity with both LOTOS and Z.

## 1 Introduction

This work is motivated by our need to specify distributed information systems and related algorithms. In particular, we require formal methods to describe a distributed database in order to provide a modelling framework for subsequent research in the areas of concurrency control, weak consistency and data replication.

To this end, considerable time has been spent in evaluating the strengths and weaknesses of two major specification languages, LOTOS [4] and Z [1]. These were the only two languages suggested to us as suitable for our purposes. Since we were mainly concerned with finding a language to use and using it rather than studying the languages themselves, we decided to limit our investigation to these two.<sup>1</sup>

The initial attraction of LOTOS is its intended use as an ISO standard language for use in OSI networking protocols. However this was also our major concern: that in targeting communication protocols as its problem domain, LOTOS might be too specific for our research needs.

Therefore we chose to compare LOTOS with Z, based on Z's growing reputation as a general purpose specification tool which has been used with much success in other research areas.

General requirements for a specification language are the ability to:

\*Research funded by Telecom Research Laboratory

<sup>1</sup>Of course, if neither language had proved suitable we would have searched further afield.

1. The underlying specifications (e.g. the Message Delivery System and Direct Naming) form an interface which defines the facilities available to the algorithm.

2. The specification of the algorithm forms an interface which is used by the distributed system. The specification of the system requires a number of definitions to be made in the specification of a process which implements the algorithm.

In the example given, the process must inherit *DirectNaming* and include *DirectNamingInitially* in the specification of its initial state. The operations *Send*, *Receive* and *InternalOperation* are provided to define interactions with the message delivery system.

In turn, the specification of the process must define the state of the process in the schema *Process*, its initial state with *ProcessInitially* and construct *ProcessOperations* as the disjunction of its operations. Additionally, it must define the type *DATA* to be equal to the disjoint union of the message types. Any constraints on the set of processes are specified in the framing schema  $\Phi$ *AlgorithmInitially*, bearing in mind that it is intended for inclusion in *DistributedSystemInitially*.

## 8 Conclusions

The reusable specification we have presented defines a distributed system of processes which communicate using a one-to-one direct-naming style of communication through some non-instantaneous medium. The specification is deliberately modular in order to reduce rewriting in future specifications.

A simple example algorithm is presented in order to demonstrate that the modules are sufficient to specify a distributed system of processes.

We intend that this document should form the beginning of a growing library of specification modules. Eventually, when specifying a new distributed algorithm, it should be possible to focus almost entirely on the details of the algorithm instead of on the surrounding system, as is currently the case.

## References

- [1] J.-R. Abrial, S. Schuman, and B. Meyer. Specification language. In R. McKeog and A. Macnaghten, editors, *On the Construction of Programs: An Advanced Course*, pages 343-410. Cambridge University Press, 1980.

- define precisely *what* is required at a high level,
- avoid prescribing *how* to do it when it is desired to allow the implementer the greatest freedom,
- successively refine the specification, possibly to the level of an implementation,
- simplify the task of verification,

We anticipate that work in the area of distributed information systems will have more specific needs, for example, the ability to:

- express the notion of processes and communication (in particular, message passing),
- specify safety and liveness requirements, e.g. termination, freedom from deadlock and fairness.

The following sections highlight significant advantages and disadvantages that we perceive in the use of LOTOS and Z for distributed information system specifications.

## 2 The Underlying Model

The most important difference between LOTOS and Z is their underlying model.

The LOTOS model is based upon concurrent *processes*, communicating through *gates*. The specification of a LOTOS process describes the behaviour sequence of that process. Processes which are *composed* together (as opposed to *interleaved*) must participate in all events (communications) occurring at their shared gates, providing synchronization points for the processes.

The Z model is based on *schemas* definitions. Some schemas define the state of an object or system (i.e. a data type), while others define legal transitions applicable to these states (i.e. the valid operations for that type). A Z specification describes the state of the system, which must be a consequence of the successive application of operations to some defined initial state.

A LOTOS specification is therefore based on a sequence of operations, while a Z specification is based on permissible states. In this sense, LOTOS gives an "active" specification, while Z provides a "passive" specification.

## 3 Fundamental Concepts of Process and Communication

Data types and functions cannot be defined in LOTOS. Instead LOTOS uses ACT-ONE (another ISO standard) for the definition of abstract data types. Although we have reservations about the use of ACT-ONE (due to its verbosity and use of equational logic), it can be used to define arbitrarily complex data (or object) types from primitive types.

Unfortunately, LOTOS introduces two new concepts, process and gate, for which there is no corresponding primitive type in ACT-ONE. Therefore processes and gates cannot be used in the definition of any other object, because LOTOS cannot define data types and ACT-ONE does not include them. By recursive process definition, it is possible to create a group of similar processes, but no mechanism exists to have a process-valued variable or parameter, nor sets or arrays of processes. Gates are declared in static lists, and again there is no facilities to describe sets or arrays of gates. The gravity of these limitations on gates is exemplified by the Topor case study specification (discussed in Section 7).

While LOTOS provides fundamental support for the notion of process and communication, Z does not. Such concepts must be specified in Z before their use. Hence we were unsure of the suitability of Z for distributed information system specifications. However once the groundwork of specifying these concepts has been performed, models of processes and communication will be types that can be used and instantiated in the same way as any other type. In [13], we present a proposal for the specification of processes and message passing in Z.

A further advantage of specifying these vital concepts in Z is that we can model processes and messages in a way most appropriate for subsequent work in distributed information systems. While LOTOS' gate-and-process model may be well-suited to OSI networking, it does not necessarily follow that it is most appropriate to distributed information systems.

## 4 Invocation of Operations and Visibility to the Environment

When specifying operations in Z, it is not clear which operations are internal to the system and which are visible to (or invoked by) the environment. The usual solution in Z is to apply a naming or scoping convention to make this distinction.

In LOTOS, only events (communications) that occur at gates declared at the outer level are visible to the environment. However there is nothing to indicate whether these events are initiated by the environment or the system. It is tempting to believe that the process performing output (sending) is the initiator, but numerous examples in the LOTOS literature suggest that this is not always so. For example, the process may be sending data in response to a request from an external process.

## 5 Specifying Sequences of Operations

A disadvantage of Z's state-based model is that it fails to provide the reader with an intuitive feel for simple operation sequences. For example, consider a simple alternating sequence of operations X and Y. The LOTOS specification captures this directly:

```
process P :=
  X ; Y ; P
endproc
```

However the Z version is certainly more cryptic with its use of pre- and post-conditions.

$$\text{WhoseTurn} \equiv \text{Xturn} | \text{Yturn}$$

```
ProcessP
turn : WhoseTurn
```

Initially

```
ProcessP
```

```
turn = Xturn
```

X

```
ΔProcessP
```

```
turn = Xturn
```

```
turn' = Yturn
```

Y

```
ΔProcessP
```

```
turn = Yturn
```

```
turn' = Xturn
```

ProcessP is a schema which defines the state of the process. The state schema *Initially* defines the initial state of ProcessP. X and Y are operation schemas, which define the only legal ways in which the state of ProcessP can change. Note that the variable *turn'* refers to value of *turn* after the operation while the undecorated *turn* refers to the value prior to the operation.

To determine that this schema is equivalent to the LOTOS specification, one must check that:

$$\begin{aligned} \text{Initially} &\Rightarrow \text{precondition}(X) \\ \text{postcondition}(X) &\Rightarrow (\text{precondition}(Y) \wedge \neg \text{precondition}(X)) \\ \text{postcondition}(Y) &\Rightarrow (\text{precondition}(X) \wedge \neg \text{precondition}(Y)) \end{aligned}$$

Some effort was devoted to trying to find some new notation to improve this aspect of Z. However it quickly became apparent that the use of comments is a more appropriate way to aid the reader's understanding, provided that the comments are used in addition to (but not as a substitute for) formal specification.

While LOTOS provides a more intuitive specification for simple operation sequencing problems, more complex and/or conditional sequences are not easy to express in the operation-sequence style in LOTOS.

## 6 Liveness Properties and Traces

Both LOTOS and Z are capable of specifying so-called "safety" properties (i.e. what must not happen) however, in order to specify liveness properties (i.e. what must happen, for example, termination), one must be able to refer to the possible sequence of events (operations) defined by the specification. For example, in a resource allocator, one might like to specify "Every requested resource must eventually be allocated".

Specifications in both LOTOS and Z define the possible sequence of events of a system (known as the set of traces). LOTOS does so by explicit sequencing and by the parallel composition of processes. In Z, the postcondition of one operation must imply the precondition of the following operation.

Although both LOTOS and Z define traces, Z specifications alone can refer to the traces within the specification itself. Since a schema is just a data type, a trace in Z is simply a sequence of "state" schemas. "Operation" schemas define the mappings from each state to the next state in the trace. Hence traces can be expressed directly in Z, and can be used to express desired liveness properties [3]. Z's ability to describe traces allows us to define temporal logic operators within Z. This is of benefit since temporal logic is an excellent notation for describing liveness properties. It has the following desirable features:

- It is a compact notation.
- It is based on a sequence of states (which corresponds exactly to a trace in Z).
- It is a formal notation, well-documented in the literature.

A temporal logic specification in Z for the resource allocation example follows with its English equivalent.

$$\begin{aligned} \square (\forall p : \text{Process} \bullet \\ p \in \text{WantResource} \Rightarrow \\ \diamond p \in \text{HasResource}) \end{aligned}$$

At all times ( $\square$ ), for all processes  $p$  which want a resource, eventually ( $\diamond$ ) each process  $p$  will have a resource.

There are many proposed temporal operators in the literature (a comprehensive set is defined in [6]), but as yet, there does not appear to be any consensus regarding the "best" set of temporal operators. However all of the proposed operators are defined over a sequence of states, and hence are amenable to formal specification in Z.

As an aside, our experiences suggest cautious use of the "next" operators in temporal logic. These operators refer to the immediate next state in which a given condition holds. In a distributed system, the overall trace of operations is an interleaving of the traces of individual processes (subject to certain restrictions). Hence two successive events in the trace of an individual process may be not be successive events in the overall trace. If "next" operators are used to express temporal requirements of the traces of individual processes, then these constraints will not necessarily hold in the overall trace and hence not be useful in proving properties at the system level.

Therefore we recommend that the use of temporal operators in distributed systems specification be restricted to such operators as  $\square$  ("always"),  $\diamond$  ("eventually") and  $\triangleleft$  (Lampert's precedence operator [7])

which apply to both the microscopic view of an individual process (or other object) and the macroscopic view of systems and subsystems.

In summary, we know of no way to express traces directly in a LOTOS specification. We can only suggest that liveness properties in LOTOS be specified in an auxiliary notation. It is possible to maintain a history of events in an individual LOTOS process, by defining this history as one of the process's parameters, and requiring that all "significant" events are recorded in this history. Such a history defines the sequence of significant events that have occurred, but cannot describe the possible events that will occur. Therefore a history mechanism (so named because it describes the past) cannot be used to specify the future. Nor is it possible to convert such "future" constraints into "past" constraints. For example, "Every resource allocation must have been previously requested" is not equivalent to "Every resource request must eventually be allocated".

A more detailed discussion of liveness constraints in LOTOS may be found in [8].

## 7 A Comparative Case Study

As a practical exercise, we "road-tested" both LOTOS and Z in the specification of an existing distributed algorithm. The algorithm chosen was a termination detection algorithm by Topor [15]. This case study example was chosen because it deals with a classic problem in distributed systems, for which we had not previously attempted a specification. Although reasonably precise and unambiguous, Topor did not attempt to specify this algorithm using any formal method (being more concerned with deriving algorithms through verification).

The algorithm assumes the existence of a connected graph with  $N$  processes (or processors) at its vertices. Only processes which are neighbours in the graph can communicate with one another, i.e. all communication takes place along the edges of the graph.

Initially we envisaged that a LOTOS gate would represent an edge of this graph (and the communication between the two processes at its endpoints). However the number of processes and the graph structure are parameters to the specification, and hence the number of such gates and the connectivity of LOTOS processes with these gates could not be accommodated by LOTOS's mechanism of declaring gates in static lists (both as parameters to processes and in the parallel composition of processes).

Reluctantly we were forced to specify all communication using a single gate. The requirement that communication occur only along the edges of the graph was enforced by a "communication medium" constraint process.

Appendix A is a high-level specification of termination detection which illustrates the solution we adopted. The final alternative in the definition of Proc is of particular significance.

Note that, as all processes share a single gate (comm) each process must participate in each I/O operation despite being neither the sender nor receiver of the communication. Each process must be specified to participate in (and ignore) any such communication at any time (or else deadlock may arise).

Conceptually, the behaviour of a process is not directly influenced by communication events between other processes, and so this solution is unnatural. Worse still, if the behaviour of the processes becomes very complex, then this method also becomes unwieldy. However given the constraint that the system contains an arbitrary number of processes connected in an arbitrary graph, no better solution is apparent. As pairwise communication within a network of processes is common in many distributed algorithms, LOTOS appears unsuited for our work.

In contrast to the unanticipated difficulty in developing a satisfactory LOTOS specification of Topor's algorithm, the development of a specification in Z progressed without incident.

Our high level specifications of termination detection in LOTOS and Z are reported in [10, 9], while the lower level specifications of Topor's algorithm are contained in [12, 11].

## 8 Software Tools

This paper does not deal with software tools in detail, being mainly concerned with the inherent expressive power of the languages. Both languages have software tools available. For example, amongst others,

LOTOS has: Galileo, Caesar, UO LOTOS, ASDE and Mana's LOTOS Compiler. Tools for Z include Fuzz, ZEN, Genesis, FORSITE and the UQ-Z editor.

The UQ-Z editor is a language-based editor for building and structuring Z specifications. The editor supports full Z syntax-checking and WYSIWYG editing with a predictive parser. This combination of features means that the user is prevented from entering syntax errors.

It is worth noting that most of the tools available for LOTOS are intended to enable the specification to be executed, whereas most of the tools for Z check syntax, types, static semantics and, in some cases, even proofs. There are exceptions to this.

ASDE [14] supports the editing and transformation of LOTOS specifications, and UO GLOTOS [2] supports similar functions for the new graphical version of LOTOS.

Researchers at British Telecom have produced ZEN [5] which helps the specifier to produce Z specifications and also supports the simulation of the specification using PROLOG.

## 9 Conclusion

While the problems which arose in the Topor case study remain unresolved, we have grave concerns about the suitability of LOTOS for distributed information systems work. Also LOTOS's apparent inability to specify traces (and hence liveness properties) within a specification further limits its value to us. Doubtless LOTOS could be altered/extended/augmented to remove these difficulties, but doing so would negate what is perhaps LOTOS's major attraction, i.e. its status as an ISO standard.

For the work we will be undertaking, Z has some disadvantages:

1. no built-in notion of process or communication,
2. simple operation sequences are not immediately obvious,
3. uncertainties regarding invocation of operations and visibility of operations to the environment.

However we believe that these drawbacks can be resolved in the following ways:

1. develop "library" schemas for processes and communication (as in [13]),
2. use comments or explanatory text to increase readability,
3. adopt a naming or scoping convention.

Therefore for specifications of distributed systems and algorithms, Z has many advantages and no insurmountable disadvantages. For these reasons, we have chosen Z as the more appropriate formal description technique for distributed information systems research.

## Acknowledgements

We thank Dr Roger Duke, Mr Graeme Smith and Dr Ian Hayes for their advice about Z, and Dr Bob Pascoe (University of the Northern Territory) and Prof Ken Turner (University of Stirling, Scotland) for their LOTOS comments.

## References

- [1] J.-R. Abrial, S. Schrumman, and B. Meyer. Specification language. In R. McKeog and A. Macnaghten, editors, *On the Construction of Programs: An Advanced Course*, pages 343-410. Cambridge University Press, 1980.
- [2] T.Y. Cheung. Uo-glotos - a system for representing, editing and translating graphical lotos. In *Proceedings of the Second International Conference on Formal Description Techniques*, 1989.

[3] Roger Duke and Graeme Smith. Temporal logic and Z specifications. In *Proceedings of 12th Australian Computer Science Conference*, Wollongong, Australia, February 1989.

[4] International Organization for Standardization, Geneva. *Information Processing Systems - Open Systems Interconnection - LOTOS, A Formal Description Technique based on the Temporal Ordering of Observation Behaviour*, July 1987.

[5] K.P. Ishaq, J.J. Masterson and A.T. Hockley. An approach to providing support tools for formal specification. In K.J. Turner, editor, *Formal Description Techniques*, pages 1-14. North-Holland, 1989.

[6] Fred Kroeger. *Temporal Logic of Programs*, volume 8 of *Monographs on Theoretical Computer Science*. Springer-Verlag, Berlin, 1987.

[7] Leslie Lamport. What good is temporal logic? *Information Processing 83*, pages 657-668, 1983.

[8] Tim Mansfield and Kerry Raymond. Expressing liveness constraints in LOTOS, March 1989. University of Queensland, Dept of Computer Science, DSL-DD10.

[9] Kerry Raymond and Tim Mansfield. A high-level specification of distributed termination in Z, January 1989. University of Queensland, Dept of Computer Science, DSL-DD3.

[10] Kerry Raymond and Tim Mansfield. High level specification of termination detection in a distributed system in LOTOS, January 1989. University of Queensland, Dept of Computer Science, DSL-DD1.

[11] Kerry Raymond and Tim Mansfield. A specification in Z of Topor's distributed termination detection algorithm, January 1989. University of Queensland, Dept of Computer Science, DSL-DD4.

[12] Kerry Raymond and Tim Mansfield. A specification in LOTOS of Topor's distributed termination detection algorithm, January 1989. University of Queensland, Dept of Computer Science, DSL-DD2.

[13] Kerry Raymond and Tim Mansfield. Specification of a distributed system, May 1989. University of Queensland, Dept of Computer Science, DSL-DD8.

[14] Saturnino Rodriguez. Asde, a tool for editing, transforming, and animating lotos specifications. In *Proceedings of the Second International Conference on Formal Description Techniques*, 1989.

[15] R.W. Topor. Termination detection for distributed computations. *Information Processing Letters*, 18:33-36, 1984.

## Appendix A

### LOTOS Specification of Termination Detection

For brevity, we have excluded the ACT-ONE type definitions. Relevant definitions are:

connections is a set of pairs. The operation neighbours checks whether a given pair is in a set of connections. status which is an enumerated type intended to record the status of a process, permitted values are active or idle.

message which is any message sent between processes.

Note that Proc are identified by natural numbers.

```

process Proc [term, comm] (me : nat, status : activity) :=
  (* Models processes, 'me' is the process ID, 'status' is
  whether the process is active or idle *)
  (* active processes may become idle *)
  [status = active] ->
  i: Proc [term, comm] ( me, idle)
□
  (* active processes may send a message *)
  [status = active] ->
  comm !me ?id: nat !message;
  Proc [term, comm] ( me, active)
□
  (* any process may receive a message, which may make it
  active *)
  comm ?id: nat !me !message;
  ( i: Proc [term, comm] ( me, status)
  □
  i: Proc [term, comm] ( me, active)
  )
□
  (* an idle process is always prepared to terminate if
  everyone else agrees *)
  [status = idle] -> term !done; exit
□
  (* every process is willing to allow pairs of processes
  of which it isn't a member, to communicate *)
  comm ?id1: nat ?id2: nat ?m: msgType [id1 != me and id2 != me];
  Proc [term, comm] (me, status)
endproc (* Proc *)

endspec (* Termination *)

```

```

specification Termination [term, comm] (numproc: nat,
  edges: connections) :=
  (* A High level description of termination detection in
  a distributed system. The system can be said to have
  terminated when all the processes in the system synchronize
  on a 'done' event at the 'term' gate. *)
  behaviour
  GenProc [term, comm] (numproc)
  ||
  CommsMedium [term, comm] (edges)
  where
  process CommsMedium [term, comm] (edges: connections) :=
  (* Permits communication only between processes connected
  by the graph defined by 'edges'. *)
  comm ?from: nat ?to: nat ?m: msgType [neighbours(edges, from, to)]
  CommsMedium [comm, term] (edges)
□
  (* terminates along with everything else *)
  term !done; exit
endproc (* CommsMedium *)

process GenProc [term, comm] (numproc : nat) :=
  (* Generates numproc processes *)
  (* All processes start 'active'. *)
  [numproc > 1] -> Proc [term, comm] (numproc, active)
  || GenProc [term, comm] (numproc - 1)
□
  [numproc = 1] -> Proc [term, comm] (numproc, active)
endproc (* GenProc *)

```