

Number 71



UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

Executing temporal logic programs (full version)

Ben Moszkowski

August 1985

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<https://www.cl.cam.ac.uk/>

© 1985 Ben Moszkowski

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

<https://www.cl.cam.ac.uk/techreports/>

ISSN 1476-2986

DOI *<https://doi.org/10.48456/tr-71>*

Executing Temporal Logic Programs* (full version)

Ben Moszkowski

Computer Laboratory, University of Cambridge,
Corn Exchange Street, Cambridge CB2 3QG, England

5 August 1985

Abstract

Over the last few years, temporal logic has been investigated as a tool for reasoning about computer programs, digital circuits and message-passing systems. In earlier work we proposed that temporal logic itself directly serve as the basis of a programming language. Since then we have implemented an interpreter for such a language called *Tempura*. We are developing Tempura as a tool for directly executing suitable temporal logic specifications of digital circuits, parallel programs and other dynamic systems. Since every Tempura statement is also a temporal formula, we can use the entire temporal logic formalism for our assertion language and semantics. Tempura has the two seemingly contradictory properties of being a logic programming language and having imperative constructs such as assignment statements.

The presentation given here describes *Interval Temporal Logic* (ITL), a formalism with conventional temporal operators such as \bigcirc (*next*) and \square (*always*) as well as lesser known ones such as *chop*. This serves as the basis for the Tempura programming language. We show the design of an interpreter for Tempura and include a variety of sample Tempura programs illustrating how to model both hardware and software. The concept of *temporal projection* is also discussed and incorporated into ITL and Tempura.

*An earlier version of this is scheduled to appear in the *Proceedings of the NSF/SERC Workshop on the Semantics of Concurrency*, Carnegie-Mellon University, Pittsburgh, Pennsylvania, USA, July 9–11, 1984.

Contents

| | |
|---|-------------|
| Table of Contents | ii |
| List of Figures | viii |
| 1 Introduction | 1 |
| 1.1 Organization of Paper | 1 |
| 2 Basic Features of the Temporal Logic | 2 |
| 2.1 Syntax of the Logic | 2 |
| 2.1.1 Syntax of expressions | 2 |
| 2.1.2 Syntax of formulas | 3 |
| 2.2 Models | 3 |
| 2.3 Interpretation of Expressions and Formulas | 4 |
| 2.4 Satisfiability and Validity | 6 |
| 3 Deriving Other Operators | 7 |
| 3.1 Boolean Operators | 7 |
| 3.2 The Operators <i>empty</i> and <i>more</i> | 8 |
| 3.3 The Operators <i>gets</i> and <i>stable</i> | 8 |
| 3.4 The Operator <i>halt</i> | 10 |
| 3.5 Temporal Equality | 10 |
| 3.6 The Operator <i>weak next</i> | 11 |
| 4 A Temporal Programming Language | 11 |
| 4.1 Syntax of Tempura | 13 |
| 4.1.1 Locations | 13 |
| 4.1.2 Expressions | 13 |

| | | |
|----------|--|-----------|
| 4.1.3 | Statements | 14 |
| 4.2 | Some Other Statements | 15 |
| 4.3 | Determinism | 15 |
| 5 | Some Additional Constructs | 16 |
| 5.1 | Static Variables | 16 |
| 5.2 | Quantified Formulas | 17 |
| 5.3 | Enlarging the Data Domain | 18 |
| 5.4 | The Operator <i>len</i> | 20 |
| 5.5 | Bounded Quantification | 20 |
| 5.6 | The Operator <i>fin</i> | 21 |
| 5.7 | Temporal Assignment | 22 |
| 5.8 | Incorporating these Constructs into Tempura | 23 |
| 5.8.1 | Locations | 23 |
| 5.8.2 | Expressions | 24 |
| 5.8.3 | Statements | 24 |
| 6 | The Operator <i>chop</i> and Related Constructs | 25 |
| 6.1 | Syntax and Semantics of <i>chop</i> | 25 |
| 6.2 | Discussion of the Operator <i>chop</i> | 26 |
| 6.3 | Simple For-Loops | 27 |
| 6.4 | Indexed For-Loops | 27 |
| 6.5 | While-Loops and Related Constructs | 28 |
| 6.6 | Deriving For-Loops and While-Loops | 29 |
| 6.7 | The Construct <i>skip</i> | 30 |
| 6.8 | Incorporating these Constructs into Tempura | 31 |

| | | |
|----------|---|-----------|
| 7 | Some Applications | 32 |
| 7.1 | Tree Summation | 32 |
| 7.1.1 | Serial solution | 34 |
| 7.1.2 | Parallel solution | 35 |
| 7.1.3 | Correctness and performance | 38 |
| 7.2 | Partitioning a List | 38 |
| 7.2.1 | Correctness of <i>partition_list</i> | 40 |
| 7.3 | Quicksort | 41 |
| 7.3.1 | Explanation of <i>quick_partition</i> | 44 |
| 7.3.2 | Explanation of <i>serial_sort_parts</i> | 45 |
| 7.3.3 | Parallel quicksort | 45 |
| 7.4 | A Multiplication Circuit | 45 |
| 7.5 | Pulse Generation | 47 |
| 7.6 | Testing a Latch | 49 |
| 7.7 | Synchronized Communication | 53 |
| 8 | Design of an Interpreter for Tempura | 55 |
| 8.1 | Variables Used by the Interpreter | 60 |
| 8.2 | Basic Execution Algorithm | 61 |
| 8.3 | Description of the Procedure <i>execute_single_state</i> | 63 |
| 8.4 | Description of the Procedure <i>transform_stmt</i> | 63 |
| 8.4.1 | Implementing the statements <i>true</i> and <i>false</i> | 64 |
| 8.4.2 | Implementing equalities | 64 |
| 8.4.3 | Implementing the statements <i>empty</i> and <i>more</i> | 65 |
| 8.4.4 | Implementing the <i>request</i> and <i>display</i> statements | 65 |
| 8.4.5 | Implementing conjunctions | 65 |

| | | |
|--------|---|----|
| 8.4.6 | Transforming the statements $\odot w$, $\circ w$ and $\square w$ | 66 |
| 8.4.7 | Implementing implications | 66 |
| 8.4.8 | Implementing some other statements | 66 |
| 8.5 | Implementing Locations | 66 |
| 8.6 | Implementing Expressions | 67 |
| 8.7 | Implementing Static Variables, Lists and Quantifiers | 67 |
| 8.7.1 | Implementing static variables | 67 |
| 8.7.2 | Implementing lists | 68 |
| 8.7.3 | Implementing the operators <i>list</i> , <i>fixed_list</i> and <i>stable_struct</i> | 68 |
| 8.7.4 | Implementing list constructors | 69 |
| 8.7.5 | Implementing subscripts | 69 |
| 8.7.6 | Implementing the list-length operator | 69 |
| 8.7.7 | Implementing existential quantifiers | 69 |
| 8.7.8 | Implementing bounded universal quantification | 70 |
| 8.7.9 | Implementing the operators <i>len</i> , <i>fin</i> and \leftarrow | 70 |
| 8.7.10 | Implementing predicate and function definitions | 71 |
| 8.7.11 | Implementing predicate and function invocations | 71 |
| 8.8 | Implementing <i>chop</i> and Iterative Operators | 72 |
| 8.8.1 | Implementing the operator <i>chop</i> | 72 |
| 8.8.2 | Implementing for-loops and other iterative statements | 73 |
| 8.8.3 | Implementing the operator <i>skip</i> | 74 |
| 8.9 | Alternative Interpreters | 74 |
| 8.9.1 | Immediate assignments | 74 |
| 8.9.2 | Two-level memory | 75 |
| 8.9.3 | Single-pass processing | 75 |
| 8.9.4 | Time stamps | 76 |

| | | |
|-----------|--|-----------|
| 8.9.5 | Ignoring the operator <i>stable</i> | 76 |
| 8.9.6 | Redundant assignments | 76 |
| 8.9.7 | Special-purpose constructs | 77 |
| 8.9.8 | Suppressing checks | 77 |
| 8.9.9 | Call-by-name | 78 |
| 9 | Experimental Features | 79 |
| 9.1 | Temporal Projection | 79 |
| 9.1.1 | Incorporating temporal projection in Tempura | 81 |
| 9.1.2 | Universal projection | 82 |
| 9.2 | Lambda Expressions and Pointers | 83 |
| 9.2.1 | Lambda expressions | 83 |
| 9.2.2 | Pointers | 84 |
| 9.3 | The <i>process</i> Construct | 85 |
| 9.4 | Implementing these Constructs | 85 |
| 9.4.1 | Implementing temporal projection | 86 |
| 9.4.2 | Implementing lambda expressions and pointers | 86 |
| 9.4.3 | Implementing the <i>process</i> construct | 87 |
| 10 | Experience and Further Work | 87 |
| 11 | Some Related Work | 88 |
| 11.1 | Other Temporal Logics Based On Intervals | 88 |
| 11.2 | The Programming Language <i>Lucid</i> | 89 |
| 11.3 | <i>CCS</i> and <i>CSP</i> | 90 |
| 11.4 | Predicative Programming | 90 |
| 11.5 | Synthesis and Verification of Specifications | 90 |
| 11.6 | The Programming Language <i>Prolog</i> | 91 |
| 11.7 | Functional Programming | 91 |

| | |
|-------------------------|-----------|
| 12 Conclusions | 92 |
| Acknowledgements | 92 |
| References | 93 |

List of Figures

| | | |
|----|--|----|
| 1 | Execution of formula (2) | 12 |
| 2 | Execution of formula (3) | 13 |
| 3 | Execution of formula (4) | 21 |
| 4 | Execution of formula (5) | 31 |
| 5 | Two binary trees | 33 |
| 6 | Predicates to sum the nodes of a binary tree in serial | 34 |
| 7 | Two executions of serial tree summation | 35 |
| 8 | Predicates to sum the nodes of a binary tree in parallel | 36 |
| 9 | Two executions of parallel tree summation | 37 |
| 10 | Predicates for partitioning a list | 39 |
| 11 | Execution of predicate for partitioning | 40 |
| 12 | Predicates for serial quicksort | 43 |
| 13 | Execution of serial quicksort | 44 |
| 14 | Predicates for parallel quicksort | 46 |
| 15 | Execution of parallel quicksort | 47 |
| 16 | Block diagram of multiplication circuit | 48 |
| 17 | Definitions of individual components | 49 |
| 18 | Definition of multiplier | 50 |
| 19 | Simulation of multiplier | 50 |
| 20 | Program to generate some digital waveforms | 51 |
| 21 | Execution of waveform generator | 51 |
| 22 | Block diagram of SR-latch | 52 |
| 23 | SR-latch operation | 52 |
| 24 | Program to simulate SR-latch | 53 |

| | | |
|----|--|----|
| 25 | Simulation of SR-latch | 54 |
| 26 | Block diagram of sender-receiver system | 55 |
| 27 | Execution of sender-receiver system | 56 |
| 28 | Block diagram of parallel lexer-parser-evaluator | 57 |
| 29 | Heart of parallel lexer-parser-evaluator | 57 |
| 30 | Execution of lexer-parser-evaluator system | 58 |
| 31 | Transformation of formula (6) | 59 |
| 32 | Details of execution of formula (6) | 62 |
| 33 | Example of projection | 80 |
| 34 | Execution of formula (7) | 81 |
| 35 | Execution of formula (8) | 82 |

1 Introduction

Temporal logic [19,30] has been recently put forward as a useful tool for reasoning about concurrent programs and hardware. Within temporal logic, one can express logical operators corresponding to time-dependent concepts such as “*always*” and “*sometimes*.” Consider, for example, the English sentence

“*If the propositions P and Q are always true, then P is always true.*”

This can be represented in temporal logic by the formula

$$\Box(P \wedge Q) \supset \Box P.$$

Here the operator \Box corresponds to the notion “*always*.” Thus, the subformula $\Box(P \wedge Q)$ can be understood as “ *P and Q are always true.*”

Typically, temporal logic has been thought of as a tool for specifying and proving properties of programs written in, say, Hoare’s CSP [15] or variants of Pascal with concurrency [14]. This distinction between temporal logic and programming languages has troubled us since it has meant that we must simultaneously use two separate notations. Programming formalisms such as Hoare logic [13], dynamic logic [9,29], and process logic [3,10] also reflect this dichotomy. One way to bridge the gap is by finding ways of using temporal logic itself as a tool for programming and simulation. With this in mind, we have developed *Tempura*, an imperative programming language based on subsets of temporal logic. Every Tempura statement is a temporal logic formula. This lets us specify and reason about Tempura programs without the need for two notations.

The underlying formalism used is called *Interval Temporal Logic* (ITL) [25,8,27] and includes such conventional temporal operators as \bigcirc (*next*) and \Box (*always*) as well as lesser known ones such as *chop*. This provides a basis for the Tempura programming language. We present ITL and Tempura and describe several sample Tempura programs illustrating how to model the structure and behavior of hardware and software systems in a unified way. The design of an interpreter for Tempura is also discussed.

1.1 Organization of Paper

We start off in section 2 by reviewing the syntax and semantics of a temporal logic having the operators \Box (*always*) and \bigcirc (*next*). In section 3 a number

of temporal constructs are then presented and then used in section 4 to build legal Tempura programs. We extend ITL and Tempura in section 5 to include constructs for list structures as well as existential and universal quantifiers. The operator *chop* is introduced in section 6 and used to express for-loops and other iterative constructs. A variety of Temporal programmings examples are then given in section 7 to show the utility of the language in dealing with hardware and software. After this is a discussion in section 8 on the details of implementing an interpreter for Tempura. We then investigate the concept of *temporal projection* as well as lambda expressions and pointers in section 9. Section 10 looks at the current status of work on Tempura interpreters. In section 11 we compare our approach with related research.

2 Basic Features of the Temporal Logic

Before describing Tempura, it is necessary to have an understanding of the underlying temporal logic. Some of the constructs described here will be later used in Tempura programs. Others will facilitate reasoning about program behavior. Rather than presenting the entire logic at once, we will first introduce some basic operators and derive others from them. In a later section, some additional operators will be considered.

2.1 Syntax of the Logic

The initial set of constructs includes conventional logical operators such as = (*equality*) and \wedge (*logical-and*). In addition, there are the two temporal operators \bigcirc (*next*) and \square (*always*).

2.1.1 Syntax of expressions

Expressions are built inductively as follows:

- Individual variables: A, B, C, \dots
- Functions: $f(e_1, \dots, e_k)$, where $k \geq 0$ and e_1, \dots, e_k are expressions. In practice, we use functions such as $+$ and *mod*. Constants such as 0 and 1 are treated as zero-place functions.

- Next: $\bigcirc e$, where e is an expression.

Here are two examples of syntactically legal expressions:

$$I + (\bigcirc J) + 1, \quad (\bigcirc I) + J - \bigcirc \bigcirc (I + \bigcirc J).$$

2.1.2 Syntax of formulas

Formulas are built inductively as follows:

- Predicates: $p(e_1, \dots, e_k)$, where $k \geq 0$ and e_1, \dots, e_k are expressions. Predicates include \leq and other basic relations.
- Equality: $e_1 = e_2$, where e_1 and e_2 are expressions.
- Logical connectives: $\neg w$ and $w_1 \wedge w_2$, where w, w_1 and w_2 are formulas.
- Next: $\bigcirc w$, where w is a formula.
- Always: $\square w$, where w is a formula.

Here are some syntactically legal ITL formulas:

$$(J = 2) \wedge \bigcirc (I = 3), \quad (\bigcirc \square [I = 3]) \wedge \neg ([\bigcirc J] = 4), \quad \bigcirc (\square [I = 3] \wedge \bigcirc \bigcirc [J = 4]).$$

Note that the operator \bigcirc can be used both for expressions (e.g., $\bigcirc J$) and for formulas (e.g., $\bigcirc (I = 3)$).

2.2 Models

A model is a triple $(\mathcal{D}, \Sigma, \mathcal{M})$ containing a data domain \mathcal{D} , a set of states Σ and an interpretation \mathcal{M} giving meaning to every function and predicate symbol. For the time being, we take the data domain \mathcal{D} to be the integers. A state is a function mapping variables to values in \mathcal{D} . We let Σ be the set of all such functions. For a state s in Σ and a variable A , we let $s[A]$ denote A 's value in s . Each k -place function symbol f has an interpretation $\mathcal{M}[f]$ which is a function mapping k elements in \mathcal{D} to a single value:

$$\mathcal{M}[f] \in (\mathcal{D}^k \rightarrow \mathcal{D}).$$

Interpretations of predicate symbols are similar but map to truth values:

$$\mathcal{M}[[p]] \in (\mathcal{D}^k \rightarrow \{true, false\}).$$

We assume that \mathcal{M} gives standard interpretations to operators such as $+$ and $<$.

The semantics given here keep the interpretations of function and predicate symbols independent of intervals. They can however be extended to allow for functions and predicates that take into account the dynamic behavior of parameters.

Using the states in Σ , we construct *intervals* of time from the set Σ^+ . An interval is thus any nonempty, finite sequence of states. If s , t and u are states in Σ , then the following are possible intervals:

$$\langle s \rangle, \quad \langle sttsus \rangle, \quad \langle tttt \rangle.$$

Note that an interval always contains at least one state.

We now introduce some basic notation for manipulating intervals. Let us use I to denote the set of all intervals. For the moment, we take I to be the set Σ^+ . Later on we will restrict I somewhat. Given an interval σ in I , we let $|\sigma|$ be the *length* of σ . Our convention is that an interval's length is the number of states *minus one*. Thus the intervals above have respective lengths 0, 5 and 3. The individual states of an interval σ are denoted by $\sigma_0, \sigma_1, \dots, \sigma_{|\sigma|}$. For instance, the following equality is true iff the variable A has the value 5 in σ 's final state:

$$\sigma_{|\sigma|}[[A]] = 5.$$

The model described here views time as being discrete and is not intended to be a realistic representation of the world around us. Nonetheless, it provides a sound basis for reasoning about many interesting dynamic phenomena involving timing-dependent and functional behavior. Furthermore, a discrete-time view of the world often corresponds to our mental model of digital systems and computer programs. In any case, we can always make the granularity of time arbitrarily fine.

2.3 Interpretation of Expressions and Formulas

We now extend the interpretation \mathcal{M} to give meaning to expressions and formulas on intervals. The construct $\mathcal{M}_\sigma[[e]]$ will be defined to equal the value in \mathcal{D} of

the expression e on the interval σ . Similarly, $\mathcal{M}_\sigma[w]$ will equal the truth value of the formula w on σ .

At first glance, the following definitions may seem somewhat arbitrary. We therefore suggest that an initial reading be rather cursory since the subsequent discussion and examples provide motivation. The definitions can then be referenced as needed.

- $\mathcal{M}_\sigma[v] = \sigma_0[v]$, where v is a variable.
Thus, a variable's value on an interval equals the variable's value in the interval's initial state.
- $\mathcal{M}_\sigma[f(e_1, \dots, e_k)] = \mathcal{M}[f](\mathcal{M}_\sigma[e_1], \dots, \mathcal{M}_\sigma[e_k])$.
The interpretation of the function symbol f is applied to the interpretations of e_1, \dots, e_k .
- $\mathcal{M}_\sigma[\bigcirc e] = \mathcal{M}_{\langle \sigma_1 \dots \sigma_{|\sigma|} \rangle}[e]$, if $|\sigma| \geq 1$.
We leave the value of $\bigcirc e$ unspecified on intervals having length 0.
- $\mathcal{M}_\sigma[p(e_1, \dots, e_k)] = \mathcal{M}[p](\mathcal{M}_\sigma[e_1], \dots, \mathcal{M}_\sigma[e_k])$.
- $\mathcal{M}_\sigma[e_1 = e_2] = \text{true}$ iff $\mathcal{M}_\sigma[e_1] = \mathcal{M}_\sigma[e_2]$.
- $\mathcal{M}_\sigma[\neg w] = \text{true}$ iff $\mathcal{M}_\sigma[w] = \text{false}$.
- $\mathcal{M}_\sigma[w_1 \wedge w_2] = \text{true}$ iff $\mathcal{M}_\sigma[w_1] = \text{true}$ and $\mathcal{M}_\sigma[w_2] = \text{true}$.
- $\mathcal{M}_\sigma[\bigcirc w] = \text{true}$ iff $|\sigma| \geq 1$ and $\mathcal{M}_{\langle \sigma_1 \dots \sigma_{|\sigma|} \rangle}[w] = \text{true}$.
- $\mathcal{M}_\sigma[\Box w] = \text{true}$ iff for all $i \leq |\sigma|$, $\mathcal{M}_{\langle \sigma_i \dots \sigma_{|\sigma|} \rangle}[w] = \text{true}$.

Examples

We will now illustrate the use of \mathcal{M} by considering the semantics of the sample temporal formulas given earlier. Let s , t and u be states in which the variables I and J have the following values:

| | I | J |
|-----|-----|-----|
| s | 1 | 2 |
| t | 3 | 4 |
| u | 3 | 1 |

The formula

$$(J = 2) \wedge \bigcirc(I = 3)$$

is true on an interval σ iff σ has length ≥ 1 , the value of J in the state σ_0 is 2 and the value of I in the state σ_1 is 3. Thus, the formula is true on the interval $\langle stu \rangle$. On the other hand, the formula is false on the interval $\langle ttu \rangle$ because J 's initial value on this interval is 4 instead of 2.

The formula

$$(\bigcirc \square[I = 3]) \wedge \neg([\bigcirc J] = 4)$$

is true on any interval σ having length ≥ 1 and in which I equals 3 in the states $\sigma_1, \dots, \sigma_{|\sigma|}$ and J does not equal 4 in σ_1 . Thus the formula is true on the interval $\langle sutut \rangle$ but is false on $\langle t \rangle$ and $\langle stutu \rangle$.

The formula

$$\bigcirc(\square[I = 3] \wedge \bigcirc \bigcirc[J = 4])$$

is true on an interval σ having length ≥ 3 and in which the variable I equals 3 in the states $\sigma_1, \dots, \sigma_{|\sigma|}$ and the variable J equals 4 in the state σ_3 . Thus this formula is true of the interval $\langle suutu \rangle$ but is false on $\langle s \rangle$ and $\langle sutuu \rangle$.

2.4 Satisfiability and Validity

A formula w is *satisfied* by an interval σ iff the meaning of w on σ equals *true*:

$$M_\sigma[w] = \text{true}.$$

This is denoted as follows:

$$\sigma \models w.$$

If all intervals satisfy w then w is *valid*, written $\models w$.

Example (Validity):

The following formula is true on an interval σ iff $|\sigma| \geq 1$, the variable I always equals 1 and in the state σ_1 , I equals 2:

$$\square(I = 1) \wedge \bigcirc(I = 2).$$

No interval can have all of these characteristics. Therefore the formula is false on all intervals and its negation is always true and hence valid:

$$\models \neg[\square(I = 1) \wedge \bigcirc(I = 2)].$$

3 Deriving Other Operators

The kinds of interval behavior one can describe with the constructs so far introduced may seem rather limited. In fact, this is not at all the case since we can develop quite a variety of derived operators. We will now present some derived operators that have proved useful in reasoning about simple computations.

3.1 Boolean Operators

The conventional boolean constructs $w_1 \vee w_2$ (logical-or), $w_1 \supset w_2$ (implication) and $w_1 \equiv w_2$ (equivalence) can be expressed in terms of \neg and \wedge . We can define logical-or as shown below:

$$w_1 \vee w_2 \equiv_{\text{def}} \neg(\neg w_1 \wedge \neg w_2).$$

We then express implication and equivalence as follows:

$$w_1 \supset w_2 \equiv_{\text{def}} \neg w_1 \vee w_2, \quad w_1 \equiv w_2 \equiv_{\text{def}} (w_1 \supset w_2) \wedge (w_2 \supset w_1).$$

The boolean constructs *true* and *false* can also be derived as can the conditional formula

$$\text{if } w_1 \text{ then } w_2 \text{ else } w_3.$$

Example (Implication):

If in an interval σ , the variable I always equals 1 and in the state σ_1 the variable J equals 2 then it follows that the expression $I + J$ equals 3 in σ_1 . This fact can be expressed by the following valid formula:

$$\models [\Box(I = 1) \wedge \bigcirc(J = 2)] \supset \bigcirc(I + J = 3).$$

Example (Equivalence):

The formula

$$\bigcirc([I = 1] \wedge [J = 2])$$

is true on an interval σ iff σ has length ≥ 1 and in the state σ_1 , the variable I has the value 1 and the variable J has the value 2. It turns out that the conjunction

$$\bigcirc(I = 1) \wedge \bigcirc(J = 2)$$

has the same meaning. The equivalence of these two formulas is expressible as follows:

$$\circ([I = 1] \wedge [J = 2]) \equiv [\circ(I = 1) \wedge \circ(J = 2)].$$

This formula is true on all intervals and is therefore valid. In general, if two formulas w_1 and w_2 have the same meaning on all intervals, then the equivalence $w_1 \equiv w_2$ is valid.

3.2 The Operators *empty* and *more*

The formula *empty* is true on an interval iff the interval has length 0:

$$\sigma \models \textit{empty} \quad \text{iff} \quad |\sigma| = 0.$$

We can define *empty* as follows:

$$\textit{empty} \equiv_{\text{def}} \neg \circ \textit{true}.$$

The formula *more* is true on an interval iff the interval has nonzero length. We can express *more* as follows:

$$\textit{more} \equiv_{\text{def}} \circ \textit{true}.$$

From these definitions it readily follows that *more* is the opposite of *empty*:

$$\models \textit{more} \equiv \neg \textit{empty}.$$

Example (Testing the length of an interval):

We can use the constructs \circ and *empty* to test the length of an interval. For example, the formula

$$\circ \circ \circ \textit{empty}$$

is true on an interval σ iff σ has length 3.

3.3 The Operators *gets* and *stable*

It is sometimes necessary to say that over time one expression e_1 equals another expression e_2 but with a one-unit delay. We use the construct e_1 *gets* e_2 to represent this and define it as follows:

$$e_1 \textit{ gets } e_2 \equiv_{\text{def}} \square (\textit{more} \supset [(\circ e_1) = e_2]).$$

The test *more* ensures that we do not “run off” the edge of the interval by erroneously attempting to examine the value of the expression e_1 in the nonexistent state $\sigma_{|\sigma|+1}$.

For instance, the formula *K gets 2K* is true on an interval σ iff the variable K is repeatedly doubled from each state to its successor:

$$\sigma \models K \text{ gets } 2K \quad \text{iff} \quad \text{for all } i < |\sigma|, \sigma_{i+1} \llbracket K \rrbracket = 2 \cdot \sigma_i \llbracket K \rrbracket.$$

The construct *stable e* is true iff the value of the expression e remains unchanged. We can readily define *stable* in terms of *gets*:

$$\text{stable } e \quad \equiv_{\text{def}} \quad e \text{ gets } e.$$

Example (Expressing an invariant condition):

The following formula is true on an interval σ in which I and J are both initially 0 and I repeatedly increases by 1 and J repeatedly increases by 2:

$$(I = 0) \wedge (J = 0) \wedge (I \text{ gets } I + 1) \wedge (J \text{ gets } J + 2).$$

In any interval for which this is true, J always equals $2I$. Below is a valid property that formalizes this:

$$\models [(I = 0) \wedge (J = 0) \wedge (I \text{ gets } I + 1) \wedge (J \text{ gets } J + 2)] \supset \square(J = 2I).$$

This shows how the operator \square can express an invariant condition.

Example (Stability):

The formula

$$(I = 1) \wedge \text{stable } I$$

is true iff I initially equals 1 and its value remains unchanged. This is the same as saying that I always equals 1. The following valid property expresses this equivalence:

$$\models [(I = 1) \wedge \text{stable } I] \equiv \square(I = 1).$$

3.4 The Operator *halt*

We can specify that a formula w becomes true only at the end of an interval σ by using the formula *halt* w :

$$\textit{halt } w \quad \equiv_{\text{def}} \quad \Box(w \equiv \textit{empty}).$$

Thus w must be false until the last state at which time w is true. For example, the formula

$$\textit{halt}(I > 100)$$

is true on σ iff the value of the variable I exceeds 100 in exactly the last state of σ .

Example (Repeatedly doubling a number):

From what we have so far presented, it can be seen that the formula

$$(I = 1) \wedge \textit{halt}(I > 100) \wedge (I \textit{ gets } 2I)$$

is true on an interval where the variable I is initially 1 and repeatedly doubles until it exceeds 100. The following valid implication states that intervals on which this formula is true will terminate upon I equalling the value 128:

$$\models [(I = 1) \wedge \textit{halt}(I > 100) \wedge (I \textit{ gets } 2I)] \supset \textit{halt}(I = 128).$$

3.5 Temporal Equality

The construct $e_1 \approx e_2$ is called *temporal equality* and is true iff the expressions e_1 and e_2 are always equal:

$$e_1 \approx e_2 \quad \equiv_{\text{def}} \quad \Box(e_1 = e_2).$$

Example (Computing factorials):

Consider the following formula for running through factorials:

$$(I = 0) \wedge (I \textit{ gets } I + 1) \wedge (J \approx I!).$$

The value of J can be seen to start at 1 and then repeatedly be multiplied by $I + 1$. This is expressed by the follow property:

$$\models [(I = 0) \wedge (I \textit{ gets } I + 1) \wedge (J \approx I!)] \supset [(J = 1) \wedge (J \textit{ gets } [I + 1] \cdot J)].$$

3.6 The Operator *weak next*

In order for the construct $\bigcirc w$ to be true on an interval σ , the length of σ must be at least 1. We therefore refer to this as *strong next*. The related construct $\bigoplus w$ is called *weak next* and is true on an interval σ if either σ has length 0 or the subformula w is true on $\langle \sigma_1 \dots \sigma_{|\sigma|} \rangle$. We can express *weak next* in terms of *strong next*:

$$\bigoplus w \equiv_{\text{def}} \text{empty} \vee \bigcirc w.$$

The operator *weak-next* provides a concise and natural way to express a construct as a conjunction of its immediate effect and future effect. Here are some examples:

$$\begin{aligned} \square w &\equiv w \wedge \bigoplus \square w, \\ \bigcirc w &\equiv \text{more} \wedge \bigoplus w, \\ e_1 \text{ gets } e_2 &\equiv [\text{more} \supset ([\bigcirc e_2] = e_1)] \wedge \bigoplus(e_1 \text{ gets } e_2), \\ \text{halt } w &\equiv (\text{empty} \equiv w) \wedge \bigoplus(\text{halt } w), \\ e_1 \approx e_2 &\equiv (e_1 = e_2) \wedge \bigoplus(e_1 \approx e_2). \end{aligned}$$

These kinds of equivalences turn out to be useful in the design of interpreters.

4 A Temporal Programming Language

Consider the formula

$$(M = 4) \wedge (N = 1) \wedge \text{halt}(M = 0) \wedge (M \text{ gets } M - 1) \wedge (N \text{ gets } 2N). \quad (1)$$

This holds true of intervals of length 4 in which M successively runs through the values 4, 3, 2, 1 and 0 and N simultaneously runs through the values 1, 2, 4, 8, and 16. Let us now explore how to automate the process of taking such a temporal formula and finding an interval satisfying it. One way to do this is to develop a procedure that analyzes the formula and determines the behavior of all free variables in every state. The result can itself be expressed as a temporal formula. For instance, here is one way to represent the result for formula (1):

$$\begin{aligned} &([\![M = 4]\!] \wedge [\![N = 1]\!]) \wedge \bigcirc([\![M = 3]\!] \wedge [\![N = 2]\!]) \wedge \bigcirc \bigcirc([\![M = 2]\!] \wedge [\![N = 4]\!]) \\ &\quad \wedge \bigcirc \bigcirc \bigcirc([\![M = 1]\!] \wedge [\![N = 8]\!]) \wedge \bigcirc \bigcirc \bigcirc \bigcirc([\![M = 0]\!] \wedge [\![N = 16]\!] \wedge \text{empty}). \end{aligned}$$

Note that this formula is logically equivalent to the original formula (1). We can view it as a kind of normal form containing state-by-state behavior of all free

| | | |
|-------------------------------|------|------|
| State 0: | M= 4 | N= 1 |
| State 1: | M= 3 | N= 2 |
| State 2: | M= 2 | N= 4 |
| State 3: | M= 1 | N= 8 |
| State 4: | M= 0 | N=16 |
| Done! Computation length = 4. | | |

Figure 1: Execution of formula (2)

variables. The process of determining such a normal form can be computerized. It is in essence a form of program execution where the original formula represents the program and the resulting normal form corresponds to the actual computation.

The general problem of finding a normal form for an arbitrary temporal formula is unsolvable. However, there are subsets of temporal logic for which the task is manageable. We have developed *Tempura*, a programming language based on one such subset. For example, formula (2) is a legal Tempura program which when executed produces the computation shown in figure 1.

$$(M = 4) \wedge (N = 1) \wedge \text{halt}(M = 0) \wedge (M \text{ gets } M - 1) \wedge (N \text{ gets } 2N) \wedge \square \text{display}(M, N). \quad (2)$$

This repeatedly outputs the values of M and N by means of the *display* construct. Note the the program's behavior is unaffected even if we change the order of the conjunction's operands. For instance, the following variant reverses them:

$$\square \text{display}(M, N) \wedge (N \text{ gets } 2N) \wedge (M \text{ gets } M - 1) \wedge \text{halt}(M = 0) \wedge (N = 1) \wedge (M = 4).$$

During the execution of the following program, the user is continually asked for the values of I by means of the *request* construct:

$$\square \text{request}(I) \wedge \text{halt}(I = 0) \wedge (J = 0) \wedge (J \text{ gets } J + I) \wedge \square \text{display}(J). \quad (3)$$

These values are summed into J and J itself is displayed. The interval terminates upon I equalling 0. A typical session is given in figure 2. Numbers in boxes (e.g., 6) are input by the user.

Let us now look at the basic syntax of Tempura. In later sections, as new temporal logic constructs are introduced variants of them will be added to Tempura.

| | |
|-------------------------------|----------|
| State 0: I = | 6 |
| State 0: J = | 0 |
| State 1: I = | 2 |
| State 1: J = | 6 |
| State 2: I = | 5 |
| State 2: J = | 8 |
| State 3: I = | 0 |
| State 3: J = | 13 |
| Done! Computation length = 3. | |

Figure 2: Execution of formula (3)

4.1 Syntax of Tempura

The main syntactic categories in Tempura are locations, expressions and statements. Let us look at each of these separately:

4.1.1 Locations

A *location* is a place into which values can be stored and later retrieved. Variables such as I , J and K are permissible locations. In addition, if l is a location, so is the temporal construct $\circ l$.

4.1.2 Expressions

Expressions can be either arithmetic or boolean. All numeric constants and variables are legal arithmetic expressions. In addition, if e_1 and e_2 are arithmetic expressions, so are operations such as the following:

$$e_1 + e_2, \quad e_1 - e_2, \quad e_1 \cdot e_2, \quad e_1 \div e_2, \quad e_1 \bmod e_2.$$

Relations such as $e_1 = e_2$ and $e_1 \geq e_2$ are boolean expressions. If b , b_1 and b_2 are boolean expressions, then so are the following:

$$\neg b, \quad b_1 \wedge b_2, \quad b_1 \vee b_2, \quad b_1 \supset b_2, \quad b_1 \equiv b_2.$$

The constants *true* and *false*, the constructs *empty* and *more* are boolean expressions as well.

The conditional expression

$$\textit{if } b \textit{ then } e_1 \textit{ else } e_2$$

is permitted. Here e_1 and e_2 can be either arithmetic or boolean expressions.

4.1.3 Statements

Certain temporal formulas are legal statements in Tempura. A statement is either *simple* or *compound*. Simple statements are built from the constructs given below. Here l is a location, e an arithmetic expression and b a boolean expression:

| | |
|--------------|---------------------|
| <i>true</i> | (no-operation) |
| <i>false</i> | (abort) |
| $l = e$ | (simple assignment) |
| <i>empty</i> | (terminate) |
| <i>more</i> | (do not terminate). |

The statement $l = e$ stores the value of the arithmetic expression e into the location l . In addition to these statements, the following can be used for requesting and displaying values:

| | |
|-------------------------------------|----------------------------------|
| $\textit{request}(l_1, \dots, l_n)$ | (request values of locations) |
| $\textit{display}(e_1, \dots, e_n)$ | (display values of expressions). |

Compound statements are built from the constructs given below. Here w , w_1 and w_2 are themselves statements and b is a boolean expression:

| | |
|------------------|------------------------|
| $w_1 \wedge w_2$ | (parallel composition) |
| $b \supset w$ | (implication) |
| $\odot w$ | (weak next) |
| $\square w$ | (always) |

Note that certain temporal formulas can be used as both boolean expressions and statements. Here are three examples:

$$I = 3, \quad (J = 2) \wedge (K = J + 3), \quad (I = 0) \supset \textit{empty}.$$

On the other hand, the following legal boolean expressions are not Tempura statements even though they are semantically equivalent to the respective formulas given above:

$$3 = I, \quad (2 = J) \wedge (J + 3 = K), \quad \neg(I = 0) \vee \textit{empty}.$$

4.2 Some Other Statements

Other constructs such as *gets*, *stable* and *halt* can be readily added to Tempura. One way to do this is to expand these to statements already described. Here are some possible equivalences:

$$\begin{aligned} \textit{if } b \textit{ then } w_1 \textit{ else } w_2 &\equiv (b \supset w_1) \wedge (\neg b \supset w_2), \\ \bigcirc w &\equiv \textit{more} \wedge \textcircled{w} w, \\ \textit{l gets e} &\equiv \square(\neg \textit{empty} \supset [(\bigcirc l) = e]), \\ \textit{stable l} &\equiv \textit{l gets l}, \\ \textit{halt b} &\equiv \square(\textit{if } b \textit{ then empty else more}), \\ l \approx e &\equiv \square(l = e). \end{aligned}$$

An alternative approach is to include these features directly in the base language.

4.3 Determinism

As we mentioned earlier, Tempura statements are limited to a subset of temporal formulas. However, even syntactically legal programs may possibly not be executable. This is because the interpreter expects the user to completely specify the behavior of program variables and to indicate when termination should occur. For example, the formula

$$I \textit{ gets } (I + 1)$$

lacks information on I 's initial value and does not specify when to stop. Thus it cannot by itself be transformed to any particular computation sequence on I and is therefore not considered a complete program. Other details must be included for the interpreter to operate properly. Of course, we could be more lenient by using backtracking and related techniques to resolve such omissions. However, for the sake of the simplicity and efficiency of the interpreter, it seems reasonable at the moment to require explicit, unambiguous information on all aspects of variable behavior.

5 Some Additional Constructs

Let us now consider how to add three important features to the temporal logic. These are *static variables*, *lists* and *quantifiers*. We subsequently use them as a basis for deriving one operator that specifies interval length and another that describes in-place assignment.

5.1 Static Variables

For a given variable A and an interval σ , it is possible for A to have a different value in each of σ 's states $\sigma_0, \sigma_1, \dots, \sigma_{|\sigma|}$. For this reason, A is called a *state variable*. It turns out to be useful to introduce a category of variables called *static variables*. Our convention is to have them start with a lower-case letter (e.g., a , $tree$ and aBc). We now constrain the set I of permitted intervals so that for any interval σ in I and any static variable a , the values of a on σ 's states are all identical:

$$\mathcal{M}_{\sigma_0}[a] = \mathcal{M}_{\sigma_1}[a] = \dots = \mathcal{M}_{\sigma_{|\sigma|}}[a].$$

Thus a static variable is stable:

$$\models \text{stable } a.$$

Identifiers starting with upper-case letters (e.g., A and $Tree$) remain state variables and can change from state to state. Furthermore, even a static variable can have different values on two distinct intervals σ and σ' .

Example (Computing powers using static variables):

The following formula has the variable J successively equal the powers $m^0, m^1, m^2, \dots, m^n$:

$$(I = 0) \wedge (J = 1) \wedge (I \text{ gets } I + 1) \wedge (J \text{ gets } m \cdot J) \wedge \text{halt}(I = n).$$

The variables m and n are static and therefore do not need to be kept explicitly stable.

5.2 Quantified Formulas

We permit formulas of the form

$$\exists V: w,$$

where V is any variable and w is itself a formula. This is called *existential quantification*. Note that V can be either a static or state variable. Below is an instance of this construct:

$$\exists I: \square(J = 2I).$$

Existential quantification readily generalizes to many variables:

$$\exists V_1, V_2, \dots, V_n: w \equiv \exists V_1: (\exists V_2: (\dots (\exists V_n: w))).$$

Universal quantification has the form $\forall V: w$ and is defined as the dual of existential quantification:

$$\forall V: w \equiv_{\text{def}} \neg \exists V: \neg w.$$

Here are the semantics of \exists :

$$\begin{aligned} \mathcal{M}_\sigma[\exists V: w] = \text{true} \quad \text{iff} \\ \text{for some interval } \sigma' \in I, \sigma \sim_V \sigma' \text{ and } \mathcal{M}_{\sigma'}[w] = \text{true}. \end{aligned}$$

The relation $\sigma \sim_V \sigma'$ is defined to be true iff the intervals σ and σ' have the same length and agree on the behavior of all variables except possibly the variable V .

Example

Consider the following states and their assignments to the variables I and J :

| | I | J |
|-----|-----|-----|
| s | 2 | 4 |
| t | 0 | 4 |
| u | 2 | 3 |

We assume that s , t and u agree on assignments to all other variables.

The formula

$$\exists I. \square(J = 2I)$$

is intuitively true on any interval on which we can construct an I such that J always equals $2I$. This is the same as saying that J is always even. For example,

the interval $\langle ttt \rangle$ satisfies the formula. From the semantics of \exists given previously it follows that to show this we need to construct an interval σ' for which the relation $\langle ttt \rangle \sim_I \sigma'$ is true and which satisfies the subformula $\Box(J = 2I)$. The interval $\langle sss \rangle$ achieves both of these constraints. Therefore $\langle ttt \rangle$ satisfies the original formula. Other intervals satisfying the formula include $\langle sss \rangle$ itself and $\langle sst \rangle$ but not $\langle u \rangle$ or $\langle stut \rangle$. Existential quantification is a tricky concept and the reader should not necessarily expect to grasp it immediately.

Example (Hiding a variable):

The formula below has J always equalling twice the value of a hidden variable I . The value of I is initially 0 and repeatedly increases by 1:

$$\exists I: [(I = 0) \wedge (I \text{ gets } I + 1) \wedge (J \approx 2I)].$$

This is logically equivalent to initializing J to 0 and repeatedly increasing it by 2:

$$(J = 0) \wedge (J \text{ gets } J + 2).$$

We can express this equivalence as the following property:

$$\models (\exists I: [(I = 0) \wedge (I \text{ gets } I + 1) \wedge (J \approx 2I)]) \equiv [(J = 0) \wedge (J \text{ gets } J + 2)].$$

The following formula has two distinct variables both called I :

$$(I = 0) \wedge (I \text{ gets } I + 1) \wedge \text{halt}(I = 5) \wedge \exists I: [(I = 1) \wedge (I \text{ gets } 3I)].$$

The first I runs from 0 to 5 and in parallel the second I is repeatedly tripled from 1 to 243. The use of existential quantification (\exists) keeps the two I 's separate and in effect hides the second one. In fact the formula is logically equivalent to its subformula

$$(I = 0) \wedge (I \text{ gets } I + 1) \wedge \text{halt}(I = 5).$$

As these examples illustrate, the operator \exists provides of means of creating locally scoped variables.

5.3 Enlarging the Data Domain

Until now we have assumed that the underlying data domain \mathcal{D} consists of the integers. Let us now enlarge it to include the boolean values *true* and *false* as well

as nested finite lists of values. Here are some sample lists:

$$\langle 3 \rangle, \langle \rangle, \langle \langle true, 2 \rangle, \langle \rangle, \langle 1, \langle 2, false \rangle \rangle \rangle.$$

The following constructs are now included among the temporal logic's expressions and are later added to Tempura:

- Simple list construction: $\langle e_0, \dots, e_{n-1} \rangle$,
where $n \geq 0$ and e_0, \dots, e_{n-1} are expressions.
- Iterative list construction: $\langle e_1 : v < e_2 \rangle$,
where e_1 and e_2 are expressions and v is a static variable.
- Subscripting: $e_1[e_2]$, where e_1 and e_2 are expressions.
- List length: $|e|$, where e is an expression.

In addition, there is the following new predicate:

- List predicate: $list(e_1, e_2)$, where e_1 and e_2 are expressions.

The semantics of the various expressions for manipulating lists are as one might expect. Our convention regarding indexing lists is to start at 0 and not at 1. For example, the following expressions are all equal:

$$3, \quad |\langle false, 1, 4 \rangle|, \quad \langle 2, 3, true \rangle[1], \quad \langle 7 - i : i < 8 \rangle[4].$$

The predicate $list(e_1, e_2)$ is true on an interval if the initial value of the expression e_1 is a list of length e_2 .

When no ambiguity arises, we express a subscripted expression such as $L[i]$ using the notation L_i .

Example (Computing lists of powers):

Here is one way to compute lists of successive powers of 0, 1, 2 and 3:

$$(I = 0) \wedge (I \text{ gets } I + 1) \wedge halt(I = 10) \wedge \square(L = \langle 0^I, 1^I, 2^I, 3^I \rangle).$$

Alternatively, the following formula can be used:

$$(I = 0) \wedge (I \text{ gets } I + 1) \wedge halt(I = 10) \\ \wedge [\square list(L, 4)] \wedge (L_0 \approx 0^I) \wedge (L_1 \approx 1^I) \wedge (L_2 \approx 2^I) \wedge (L_3 \approx 3^I).$$

We now present a number of useful operators derived from the constructs so far introduced.

5.4 The Operator *len*

The formula $len(e)$ is true on an interval having length exactly e :

$$\mathcal{M}_\sigma \llbracket len(e) \rrbracket = true \quad \text{iff} \quad \mathcal{M}_\sigma \llbracket e \rrbracket = |\sigma|.$$

It turns out that we can express len by means of existential quantification and other previously introduced constructs. Here is one way:

$$len(e) \quad \equiv_{\text{def}} \quad \exists I: [(I = e) \wedge (I \text{ gets } I - 1) \wedge halt(I = 0)],$$

where I does not occur freely in e . We use I as a hidden counter that is initialized to e 's value and then keeps track of how much time remains in the interval.

Example (Doubling a variable):

The following formula has the variable N run through the first few powers of 2:

$$len(5) \wedge (N = 1) \wedge (N \text{ gets } 2N).$$

The len construct is used to specify the length of the computation.

5.5 Bounded Quantification

We now introduce formulas of the form

$$\forall v < e: w,$$

where v is a static variable, e is an expression and w is a formula. The construct is referred to as *bounded universal quantification* and can be defined as follows:

$$\forall v < e: w \quad \equiv_{\text{def}} \quad \forall v: (0 \leq v < e \supset w).$$

This is especially useful for processing elements of a list variable in parallel. Bounded existential quantification can be defined in an analogous way. In addition, it is easy to generalize this notation to handle ranges such as $v \leq e$ and $e_1 \leq v < e_2$.

| | |
|----------|-------------------------|
| State 0: | L=<1,1,1,1,1> |
| State 1: | L=<0,1,2,3,4> |
| State 2: | L=<0,1,4,9,16> |
| State 3: | L=<0,1,8,27,64> |
| State 4: | L=<0,1,16,81,256> |
| State 5: | L=<0,1,32,243,1024> |
| State 6: | L=<0,1,64,729,4096> |
| Done! | Computation length = 6. |

Figure 3: Execution of formula (4)

Example (Generalized computation of powers):

The following formula successively assigns the list L the first $n + 1$ powers of the numbers $0, 1, \dots, m - 1$:

$$\exists I: [(I = 0) \wedge (I \text{ gets } I - 1) \wedge \text{halt}(I = n) \wedge (L \approx \langle k^I : k < m \rangle)]. \quad (4)$$

Figure 3 shows the behavior of L for $m = 5$ and $n = 6$. The actual program consists of the conjunction of the above formula with the statement

$$\square \text{display}(L).$$

Here is another technique which is logically equivalent and uses bounded universal quantification to simultaneously manipulate each of the elements of L :

$$\text{len}(n) \wedge [\square \text{list}(L, m)] \wedge \forall k < m: [(L_k = 1) \wedge (L_k \text{ gets } k \cdot L_k)].$$

5.6 The Operator fin

The formula $fin w$ is true on an interval σ iff the formula w is itself true on the final subinterval $\langle \sigma_{|\sigma|} \rangle$. We express $fin w$ as follows:

$$fin w \quad \equiv_{\text{def}} \quad \square(\text{empty} \supset w).$$

The formula $fin w$ is weaker than $halt w$ since $fin w$ only looks at the last state whereas $halt w$ tests behavior throughout.

Example (Doubling a variable):

The following formula is true on an interval σ iff $|\sigma| = 3$ and I is initially 1 and repeatedly doubles:

$$\text{len}(3) \wedge (I = 1) \wedge (I \text{ gets } 2I).$$

One effect is that I ends up equal to 8. This is expressed by the valid implication given below:

$$\models [\text{len}(3) \wedge (I = 1) \wedge (I \text{ gets } 2I)] \supset \text{fin}(I = 8).$$

5.7 Temporal Assignment

The formula $e_1 \rightarrow e_2$ is true for an interval if the initial value of the expression e_1 equals the final value of the expression e_2 . We define this as follows:

$$e_1 \rightarrow e_2 \quad \equiv_{\text{def}} \quad \exists a: [(a = e_1) \wedge \text{fin}(e_2 = a)],$$

where the static variable a does not occur freely in either e_1 or e_2 . The stability of the value of a is used to compare the values of e_1 and e_2 at different times. We call this construct *temporal assignment*. For example, the formula $I + 1 \rightarrow I$ is true on an interval σ iff the value of $I + 1$ in the initial state σ_0 equals the value of I in the final state $\sigma_{|\sigma|}$. If desired, we can reverse the direction of the arrow:

$$I \leftarrow I + 1.$$

The formula

$$(I \leftarrow I + 1) \wedge (J \leftarrow J + I)$$

is then true in an interval iff I increases by 1 and in parallel J increases by I . Similarly, the following specifies that the values of the state variables A and B are exchanged:

$$(A \rightarrow B) \wedge (B \rightarrow A).$$

Unlike assignment in conventional programming languages, temporal assignment only affects variables explicitly mentioned; the values of other variables do not necessarily remain fixed. For example, the formulas

$$I \leftarrow I + 1$$

and

$$(I \leftarrow I + 1) \wedge (J \leftarrow J)$$

are not equivalent since the first formula does not require J 's initial and final values to be equal. Thus, temporal assignment lacks the so-called *frame assumption*.

Example (In-place computation of the maximum of two numbers):

The temporal formula

$$\text{if } I \geq J \text{ then } (I \leftarrow I) \text{ else } (I \leftarrow J)$$

is true in any interval where I 's value in the final state equals the maximum of the values of I and J in the initial state. This can be seen by case analysis on the test $I \geq J$.

Let the function $\text{max}(i, j)$ equal the maximum of the two values i and j . The following temporal formula therefore places the maximum of I and J into I :

$$I \leftarrow \text{max}(I, J)$$

The equivalence of the two approaches is expressed by the following property:

$$\models [I \leftarrow \text{max}(I, J)] \equiv [\text{if } I \geq J \text{ then } (I \leftarrow I) \text{ else } (I \leftarrow J)].$$

5.8 Incorporating these Constructs into Tempura

We now extend Tempura's locations, expressions and statements in the ways described below.

5.8.1 Locations

Static variables such as a and x are now permitted as locations. In addition if l is a location and e is an expression, then the subscript construct $l[e]$ is a permitted location.

5.8.2 Expressions

The following are now legal expressions:

| | |
|---------------------------------------|------------------------------|
| v | (static variables) |
| $\langle e_0, \dots, e_{n-1} \rangle$ | (simple list constructor) |
| $\langle e_1: v <, e_2 \rangle$ | (iterative list constructor) |
| $e_1[e_2]$ | (subscripting) |
| $ e $ | (list length). |

Here v is a static variable and $e, e_0, e_1, \dots, e_{n-1}$ are all expressions.

In addition, we permit function invocations of the form

$$f(e_1, \dots, e_n)$$

where f is a function defined in the manner described later and $n \geq 0$ and e_1, \dots, e_n are expressions.

5.8.3 Statements

We now allow variables to be assigned not only integers, but boolean and list values are well. Here are some examples:

$$\text{Done} = \text{true}, \quad L = \langle 1, 2, 4 \rangle, \quad \text{Flag gets } (\neg \text{Flag}).$$

In addition, the following are permitted statements:

| | |
|------------------------------|------------------------------------|
| $\text{list}(l, e)$ | (list declaration) |
| $\text{len}(e)$ | (interval length) |
| $\exists V_1, \dots, V_n: w$ | (existential quantification) |
| $\forall v < e: w$ | (bounded universal quantification) |
| $\text{fin } w$ | (terminal statement) |
| $l \leftarrow e$ | (temporal assignment). |

Here l is a location, e is an expression, V is any variable, v is a static variable, and w is a statement.

We also permit predicate invocations of the form

$$p(e_1, \dots, e_n)$$

where p is a predicate defined in the manner described below and $n \geq 0$ and e_1, \dots, e_n are expressions. The interpreter uses call-by-reference when passing parameters.

Statements of the following forms can be used to define functions and predicates respectively:

$$\begin{aligned} \text{define } f(V_1, \dots, V_n) &= e, \\ \text{define } p(V_1, \dots, V_n) &\equiv w. \end{aligned}$$

The formal parameters V_1, \dots, V_n are state variables or static variables and e is an expression and w is a statement. The identifier used in place of f or p should be static. Here are two sample definitions:

$$\begin{aligned} \text{define } \min(i, j) &= (\text{if } i \leq j \text{ then } i \text{ else } j), \\ \text{define } \text{double}(M) &\equiv (M \text{ gets } 2M). \end{aligned}$$

Note that recursive definitions are permitted. Furthermore, the body of a definition can include temporal constructs. Thus our actual usage of predicates and functions is more general than indicated in the temporal logic semantics presented earlier in section 2.2. It is not difficult to adjust the semantics to take this into account although we omit the details.

6 The Operator *chop* and Related Constructs

Temporal logic contains various constructs such as *chop* and *while* that are rather similar to certain kinds of statements found in Algol and related programming languages. We first extend the syntax and semantics of the temporal logic to include *chop*. The resulting formalism is called Interval Temporal Logic. Within it we define a number of interval-dependent operators and subsequently expand Tempura to include them.

6.1 Syntax and Semantics of *chop*

We now permit formulas of the following form:

$$w_1; w_2, \quad \text{where } w_1 \text{ and } w_2 \text{ are formulas.}$$

This operator is known as *chop*. Such a formula $w_1; w_2$ is true on an interval σ iff there is at least one way to split σ into two subintervals such that the formula w_1 is true on the left subinterval and the formula w_2 is true on the right subinterval:

$$\begin{aligned} \mathcal{M}_\sigma[w_1; w_2] = true \quad \text{iff} \\ \text{for some } i \leq |\sigma|, \mathcal{M}_{\langle \sigma_0 \dots \sigma_i \rangle}[w_1] = true \text{ and } \mathcal{M}_{\langle \sigma_i \dots \sigma_{|\sigma|} \rangle}[w_2] = true. \end{aligned}$$

Note that the two subintervals $\langle \sigma_0 \dots \sigma_i \rangle$ and $\langle \sigma_i \dots \sigma_{|\sigma|} \rangle$ share the state σ_i .

Example (Sequential composition of assignments):

The formula

$$(K + 1 \rightarrow K); (K + 2 \rightarrow K)$$

is true on an interval σ iff there is some $i \leq |\sigma|$ such that the subformula $K + 1 \rightarrow K$ is true on the subinterval $\langle \sigma_0 \dots \sigma_i \rangle$ and the subformula $K + 2 \rightarrow K$ is true on remaining subinterval $\langle \sigma_i \dots \sigma_{|\sigma|} \rangle$. The net effect is that K increases by 3. This is expressed by the following property:

$$\models [(K + 2 \rightarrow K); (K + 1 \rightarrow K)] \supset (K + 3 \rightarrow K).$$

6.2 Discussion of the Operator *chop*

The construct *chop* is rather different from the conventional temporal operators \square and \bigcirc . The latter examine an interval's suffix subintervals whereas *chop* splits the interval and tests both parts. This facilitates looking at arbitrary subintervals of time.

Harel, Kozen and Parikh [10] appear to be the first to mention *chop* as a temporal construct. It is considered in more detail by Chandra, Halpern, Meyer and Parikh [3]. In references [8] and [25] we use *chop* to facilitate reasoning about time-dependent digital hardware. Our subsequent work in [28] and [26] uses *chop* to give specifications and properties of simple algorithms and message-passing systems. In the rest of this section we examine *chop* and other ITL constructs and then extend Tempura to include them.

6.3 Simple For-Loops

The following simple for-loop repeats w for e times in succession, where w is a formula and e is an expression:

for e times do w.

This can also be written w^e . Here is one way to state that the variable I increases by 1 for 4 times in succession:

for 4 times do (I + 1 → I).

This is equivalent to the formula

$(I + 1 \rightarrow I); (I + 1 \rightarrow I); (I + 1 \rightarrow I); (I + 1 \rightarrow I).$

Since *chop* is associative, we can omit parentheses without being ambiguous. In the case of zero iterations, the for-loop is equivalent to *empty*:

\models (*for 0 times do w*) \equiv *empty*.

The formula w^* (read “*w star*”) is true if the subformula w occurs some number of times in succession. This operator is sometimes known as *chop-star*. We can express it as in terms of a simple for-loop:

$w^* \equiv_{\text{def}} \exists n: (n \geq 0 \wedge [\textit{for } n \textit{ times do } w]),$

where n does not occur freely in w . For example, the formula

$(I + 1 \rightarrow I)^*$

is true on an interval if I repeatedly increases over some unspecified number of iterations.

6.4 Indexed For-Loops

In addition to the simple for-loop, we permit an indexed variant having the following syntax:

for v < e do w.

Here v is a static variable, e is an expression and w is a formula. For instance, the formula

for k < 4 do (I + k → I).

is equivalent to

$(I + 0 \rightarrow I); (I + 1 \rightarrow I); (I + 2 \rightarrow I); (I + 3 \rightarrow I).$

6.5 While-Loops and Related Constructs

The *temporal while-loop* is another important iterative construct in ITL. The basic form is similar to that of a while-loop in Algol:

$$\textit{while } w_1 \textit{ do } w_2.$$

Both w_1 and w_2 are themselves formulas. The while-loop obeys the following general expansion property:

$$\textit{while } w_1 \textit{ do } w_2 \equiv \textit{if } w_1 \textit{ then } (w_2; [\textit{while } w_1 \textit{ do } w_2]) \textit{ else empty.}$$

Thus, if w_1 is true, the body of the loop, w_2 , is examined after which the loop is repeated. If w_1 is false, the interval must have length 0. The *chop-star* operator can in fact be derived from a while-loop:

$$w^* \equiv \textit{while more do } w.$$

A *repeat-loop* has the form

$$\textit{repeat } w_1 \textit{ until } w_2$$

and can be expressed using a while-loop:

$$\textit{repeat } w_1 \textit{ until } w_2 \equiv_{\text{def}} w_1; (\textit{while } \neg w_2 \textit{ do } w_1).$$

Another loop construct has the form

$$\textit{loop } w_1 \textit{ exit when } w_2 \textit{ otherwise } w_3,$$

where w_1 , w_2 and w_3 are themselves formulas. The formula w_2 is used to exit from the loop. Here is how to express the loop in terms of the *while* construct:

$$\textit{loop } w_1 \textit{ exit when } w_2 \textit{ otherwise } w_3 \equiv_{\text{def}} w_1; (\textit{while } \neg w_2 \textit{ do } [w_3; w_1]).$$

In section 8.2 we use this style of loop to describe the basic algorithm used by a Tempura interpreter.

Example (Computing the greatest common divisor):

Consider the following assignment which specifies that N 's final value equals the initial value of the greatest common divisor of M and N :

$$N \leftarrow \textit{gcd}(M, N).$$

The formula below implies this:

$$\textit{while } (M \neq 0) \textit{ do } ([M \leftarrow N \bmod M] \wedge [N \leftarrow M]).$$

6.6 Deriving For-Loops and While-Loops

Let us look at one way to express for-loops and while-loops entirely in terms of *chop* and other ITL constructs already introduced. This discussion can be omitted by the reader. We first define the predicate $end_points(l, n)$ to be true on intervals where l is a list of $n + 1$ offsets indicating the end points of some successive subintervals corresponding to iterations:

$$\begin{aligned} end_points(l, n) &\equiv_{\text{def}} \\ &list(l, n + 1) \wedge (l_0 = 0) \wedge len(l_n) \wedge \forall i < n: (l_i \leq l_{i+1}). \end{aligned}$$

Thus the element l_0 equals 0 since this is the start of any initial iteration. The final element l_n analogously equals the length of the interval. Furthermore, the sequence formed by l 's elements is weakly increasing.

We also make use of the construct $iteration(l, n, i, w)$. This is true on an interval if the formula w is true on the subinterval corresponding to the i -th iteration with respect to the values of l and n (i.e., the subinterval bound by l_i and l_{i+1}). Here is one way to express $iteration$:

$$iteration(l, n, i, w) \equiv_{\text{def}} len(l_i); w; len(l_n - l_{i+1}).$$

The simple for-loop construct can now be expressed as follows:

$$\begin{aligned} \text{for } e \text{ times do } w &\equiv_{\text{def}} \\ &\exists l, n: ([n = e] \wedge end_points(l, n) \wedge \forall i < n: iteration(l, n, i, w)), \end{aligned}$$

where i , l and n do not occur freely in e or w . The technique for the indexed for-loop is analogous:

$$\begin{aligned} \text{for } v < e \text{ do } w &\equiv_{\text{def}} \\ &\exists l, n: ([n = e] \wedge end_points(l, n) \wedge \forall i < n: iteration(l, n, i, \exists v: [v = i \wedge w])), \end{aligned}$$

where i , l and n do not occur freely in v , e or w .

Below is the derivation of the while-loop:

$$\begin{aligned} \text{while } w_1 \text{ do } w_2 &\equiv_{\text{def}} \\ &\exists l, n: (end_points(l, n) \\ &\quad \wedge \forall j < n + 1: [(j < n) \equiv (len(l_j); w_1)] \\ &\quad \wedge \forall i < n: iteration(l, n, i, w_2) \\ &\quad). \end{aligned}$$

Here the formula w_1 must be true at the beginning of every iteration but false at the end. Furthermore, the formula w_2 must be true on each iteration.

6.7 The Construct *skip*

The construct *skip* is true on an interval σ iff σ has length 1. We can express *skip* as follows:

$$\mathit{skip} \equiv_{\text{def}} \bigcirc \mathit{empty}.$$

Example (Measuring the length of an interval):

An interval's length can be tested using *skip* and *chop*. For example, the formula

$$\mathit{skip}; \mathit{skip}; \mathit{skip}$$

is true on intervals having length 3. It follows that this formula is equivalent to $\mathit{len}(3)$:

$$\models \mathit{len}(3) \equiv (\mathit{skip}; \mathit{skip}; \mathit{skip}).$$

Example (Unit-length iterations):

The following while-loop repeatedly decrements I and sums I into J over each unit of time until I equals 0:

$$\mathit{while} (I \neq 0) \mathit{do} (\mathit{skip} \wedge [I \leftarrow I - 1] \wedge [J \leftarrow J + I]).$$

The body of the loop contains the *skip* operator in order that the length of each iterative step be 1. The behavior can also be expressed using *halt* and *gets*. Here is a semantically equivalent way of doing this:

$$\mathit{halt}(I = 0) \wedge (I \mathit{gets} I - 1) \wedge (J \mathit{gets} J + I).$$

Example (Expressing the operator gets using skip and chop-star):

A formula such as $A \mathit{gets} B$ can be alternatively expressed as an unspecified number of iterations, each of unit length:

$$\models (A \mathit{gets} B) \equiv (\mathit{skip} \wedge [A \leftarrow B])^*.$$

| | | |
|-------------------------------|------|------|
| State 0: | I= 4 | J= 0 |
| State 1: | I= 3 | J= 4 |
| State 2: | I= 2 | J= 7 |
| State 3: | I= 1 | J= 9 |
| State 4: | I= 0 | J=10 |
| Done! Computation length = 4. | | |

Figure 4: Execution of formula (5)

6.8 Incorporating these Constructs into Tempura

We now extend Tempura to include statements based on the ITL constructs just introduce. The set of statements is expanded to include the following:

| | |
|---|--------------------------|
| <i>skip</i> | (unit interval) |
| $w_1; w_2$ | (sequential composition) |
| <i>for e times do w</i> | (simple for-loop) |
| <i>for v < e do w</i> | (indexed for-loop) |
| <i>while b do w</i> | (while-loop) |
| <i>repeat b until w</i> | (repeat-loop) |
| <i>loop w₁ exit when b otherwise w₂</i> | (exit-loop) |

Here w , w_1 and w_2 are themselves statements, e is an expression, v is a static variable and b is a boolean test.

Example (Computing sums):

The following Tempura program uses a while-loop to compute a sum:

$$(I = 4) \wedge (J = 0) \wedge \square display(I, J) \wedge [while I \neq 0 do (skip \wedge [I \leftarrow I - 1] \wedge [J \leftarrow J + I])]. \quad (5)$$

Figure 4 shows the program's behavior when run.

Example (Computing powers):

Consider the problem of finding the value of the expression I^J and placing this in another variable K . We can specify this using the temporal assignment

$$K \leftarrow I^J.$$

The following Tempura statement achieves this by looking at J 's binary structure:

$$(K = 1) \wedge [\textit{while} (J > 0) \textit{do} (\textit{skip} \wedge w)],$$

where the statement w has the form

$$\begin{aligned} &\textit{if} (J \bmod 2 = 0) \\ &\quad \textit{then} [(I \leftarrow I \cdot I) \wedge (J \leftarrow J \div 2) \wedge (K \leftarrow K)] \\ &\quad \textit{else} [(I \leftarrow I) \wedge (J \leftarrow J - 1) \wedge (K \leftarrow K \cdot I)]. \end{aligned}$$

7 Some Applications

We now present some sample Tempura programs for summing the leaves of a tree, partitioning a list and sorting a list. Afterwards programs are given for simulating a simple multiplication circuit, generating digital pulses and testing a latch. The final two examples deal with synchronized communication between parallel processes. Most of the programming constructs used here have already been introduced. Those that have not been are briefly described where first mentioned.

7.1 Tree Summation

Suppose we have a binary tree of values such as either of the ones shown in figure 5. They can be linearly represented by the following list structures:

$$\langle\langle\langle 1, 1 \rangle, \langle 1, 1 \rangle\rangle, \langle\langle 1, 1 \rangle, \langle 1, 1 \rangle\rangle\rangle, \quad \langle\langle 1, \langle 2, 3 \rangle \rangle, \langle 4, 5 \rangle\rangle.$$

Let the function $\textit{leaf_sum}(\textit{tree})$ determine the sum of a tree's leaves:

$$\begin{aligned} \textit{leaf_sum}(\textit{tree}) &=_{\textit{def}} \\ &\quad \textit{if} \textit{is_integer}(\textit{tree}) \textit{then} \textit{tree} \\ &\quad \textit{else} \textit{leaf_sum}(\textit{tree}_0) + \textit{leaf_sum}(\textit{tree}_1). \end{aligned}$$

Here the predicate $\textit{is_integer}(\textit{tree})$ is true when the parameter \textit{tree} is an integer (i.e., a leaf) and false when \textit{tree} is a pair.




Figure 5: Two binary trees

```

define serial_sum_tree(Tree) ≡
  if is_integer(Tree) then empty
  else (
    sum_subtree(Tree, 0);
    sum_subtree(Tree, 1);
    (
      skip ∧ (Tree ← Tree0 + Tree1)
    )
  ).

define sum_subtree(Tree, i) ≡
  serial_sum_tree(Treei)
  ∧ stable_struct(Tree)
  ∧ stable Tree1-i.

```

Figure 6: Predicates to sum the nodes of a binary tree in serial

Now consider the task of designing an algorithm to reduce a tree in-place to a single value indicated by *leaf_sum*. If the variable *Tree* initially equals such a binary tree, we can specify the problem as follows:

$$Tree \leftarrow leaf_sum(Tree).$$

Let us look at a serial and a parallel implementation of this.

7.1.1 Serial solution

One approach to achieving our goal is given by the predicate *serial_sum_tree* defined in figure 6. This terminates if the tree is already an integer-valued leaf. Otherwise, the predicate *sum_subtree* is used to first reduce the left subtree and then the right subtree. Finally, the statement

$$skip \wedge (Tree \leftarrow Tree_0 + Tree_1)$$

is used to reduce the tree to a single value. Note that when either of the two subtrees is being reduced by *sum_subtree*, the other one is kept stable. In addition, the built-in predicate *stable_struct* is used in the predicate *sum_subtree* to maintain the tree's main node intact so that the two subtrees can continue to be properly

```

State 0: Tree=<<<1,1>,<1,1>>,<<1,1>,<1,1>>>
State 1: Tree=<<2,<1,1>>,<<1,1>,<1,1>>>
State 2: Tree=<<2,2>,<<1,1>,<1,1>>>
State 3: Tree=<4,<<1,1>,<1,1>>>
State 4: Tree=<4,<2,<1,1>>>
State 5: Tree=<4,<2,2>>
State 6: Tree=<4,4>
State 7: Tree=8

Done! Computation length = 7.

State 0: Tree=<<1,<2,3>>,<4,5>>
State 1: Tree=<<1,5>,<4,5>>
State 2: Tree=<6,<4,5>>
State 3: Tree=<6,9>
State 4: Tree=15

Done! Computation length = 4.

```

Figure 7: Two executions of serial tree summation

accessed. We can express *stable_struct* as follows:

$$stable_struct(e) \equiv_{\text{def}} \square [more \supset list([\circ e], |e|)].$$

Thus, if the expression e is initially a list of some particular length, it remains a list of that length throughout the interval. If this operation were omitted, there would be no mention of whether the root of the tree remains a pair.

We initialize the tree and invoke *serial_sum_tree* by means of the following sort of formula:

$$\begin{aligned}
Tree = & \langle \langle \langle 1, 1 \rangle, \langle 1, 1 \rangle \rangle, \langle \langle 1, 1 \rangle, \langle 1, 1 \rangle \rangle \rangle \\
& \wedge serial_sum_tree(Tree) \wedge \square display(Tree).
\end{aligned}$$

Figure 7 shows the resulting behavior for the two different trees mentioned.

7.1.2 Parallel solution

The predicate *par_sum_tree* in figure 8 is similar to *serial_sum_tree* except that it recursively reduces each half of a pair in parallel rather than sequentially.

```

define par_sum_tree(Tree) ≡
  if is_integer(Tree) then empty
  else (
    ∃Done:(
      (Done ≈ [is_integer(Tree0) ∧ is_integer(Tree1)])
      ∧ halt(Done)
      ∧ stable_type(Tree)
      ∧ sum_tree_process(Done, Tree0)
      ∧ sum_tree_process(Done, Tree1)
    );
    (
      skip ∧ (Tree ← Tree0 + Tree1)
    )
  ).

define sum_tree_process(Done, Tree) ≡
  process(
    par_sum_tree(Tree);
    (
      halt Done
      ∧ stable Tree
    )
  ).

```

Figure 8: Predicates to sum the nodes of a binary tree in parallel

```

State 0: Tree=<<<1,1>,<1,1>>,<<1,1>,<1,1>>
State 1: Tree=<<2,2>,<2,2>>
State 2: Tree=<4,4>
State 3: Tree=8

Done! Computation length = 3.

State 0: Tree=<<1,<2,3>>,<4,5>>
State 1: Tree=<<1,5>,<9>>
State 2: Tree=<6,9>
State 3: Tree=15

Done! Computation length = 3.

```

Figure 9: Two executions of parallel tree summation

A variable named *Done* is used to monitor the progress of the two subtrees. It equals *true* when they are both finally integers. At this time the sum of the two values can be assigned to the tree variable. The subordinate predicate *sum_tree_process(Done, Tree)* reduces its tree parameter and then keeps the tree stable until the flag *Done* becomes true. This ensures that the two parallel invocations of *sum_tree_process* finish at the same time. Figure 9 shows the behavior of *par_sum_tree* on the two trees discussed above. As might be expected, the computation length is less than that required by the serial algorithm.

Note that the body of the predicate *sum_tree_process* in figure 8 is a statement of the form *process w*. The *process* construct has no special logical semantics:

$$\textit{process } w \quad \equiv_{\text{def}} \quad w.$$

It is used when several Tempura statements are run in parallel and each independently determines the interval length. See section 9.3 for a discussion.

7.1.3 Correctness and performance

ITL can be used to specify the correctness and relative speeds of the two algorithms just introduced. Here are the basic correctness properties:

$$\begin{aligned} &\models \text{serial_sum_tree}(Tree) \supset (Tree \leftarrow \text{leaf_sum}(Tree)), \\ &\models \text{par_sum_tree}(Tree) \supset (Tree \leftarrow \text{leaf_sum}(Tree)). \end{aligned}$$

The invariant and rate of progress for the serial version are shown below:

$$\begin{aligned} &\models \text{serial_sum_tree}(Tree) \supset \\ &\quad (\text{stable leaf_sum}(Tree) \\ &\quad \wedge [\text{leaf_count}(Tree) \text{ gets } \text{leaf_count}(Tree) - 1]). \end{aligned}$$

Here the function *leaf_count* equals the number of leaves in a tree. The property states that even though the tree is changing, the sum of its leaves remains stable. Furthermore, the number of leaves decreases by 1 every unit of time. The invariant for the parallel algorithm is identical to that of the serial algorithm although the rate of progress is more complicated. The tree's leaf count decreases by the number of internal nodes whose two sons are both integers.

The computation length for the serial and parallel algorithms is expressed as follows:

$$\begin{aligned} &\models \text{serial_sum_tree}(Tree) \supset \text{len}(\text{leaf_count}(Tree) - 1), \\ &\models \text{par_sum_tree}(Tree) \supset \text{len}(\text{tree_height}(Tree)). \end{aligned}$$

Here the expression $\text{leaf_count}(Tree) - 1$ equals the number of internal nodes in the tree. The function *tree_height* is defined to equal the length of the longest path from the tree's root to a leaf:

$$\begin{aligned} \text{tree_height}(tree) &=_{\text{def}} \\ &\text{if } \text{is_integer}(tree) \text{ then } 0 \\ &\text{else } \max(\text{tree_height}(tree_0), \text{tree_height}(tree_1)) \end{aligned}$$

7.2 Partitioning a List

We now describe a technique for partitioning a list. This will be subsequently used in some quicksort algorithms. The predicate

$$\text{partition_list}(L, \text{key}, \text{left_len})$$

```

define partition_list(L, key, left_len) ≡
  ∃I, J : (
    (I = 0) ∧ (J = |L|)
    ∧ (
      part_loop(L, key, I, J);
      [empty ∧ (left_len = I)]
    )
  ).

define part_loop(L, key, I, J) ≡
  while I < J do (
    skip ∧ part_step(L, key, I, J)
  )

define part_step(L, key, I, J) ≡
  if LI < key
  then (
    stable L
    ∧ (I ← I + 1) ∧ (J ← J)
  )
  else (
    swap_list(L, I, J - 1)
    ∧ (I ← I) ∧ (J ← J - 1)
  ).

```

Figure 10: Predicates for partitioning a list

| | |
|-------------------------------|-------------------|
| State 0: | L=<1,3,2,3,0,1,3> |
| State 1: | L=<1,3,2,3,0,1,3> |
| State 2: | L=<1,3,2,3,0,1,3> |
| State 3: | L=<1,1,2,3,0,3,3> |
| State 4: | L=<1,1,2,3,0,3,3> |
| State 5: | L=<1,1,0,3,2,3,3> |
| State 6: | L=<1,1,0,3,2,3,3> |
| State 7: | L=<1,1,0,3,2,3,3> |
| State 7: | left_len= 3 |
| Done! Computation length = 7. | |

Figure 11: Execution of predicate for partitioning

defined in figure 10 uses the predicate *part_loop* to iteratively permute the elements of the list parameter *L* so that those less than the value of *key* end up to the left of elements greater than or equal to *key*. Each step of the loop invokes the predicate *part_step*. The length of the left side is finally stored in the static parameter *left_len*. For example, suppose that *L* initially equals the list $\langle 1, 3, 2, 3, 0, 1, 3 \rangle$ and *key* has the value 2. The final value of *L* is then the list $\langle 1, 1, 0, 3, 2, 3, 3 \rangle$ and the value of *left_len* is 3. In figure 11 we depict the behavior of *L* in each state and display the value of *left_len* in the final state. Note that the predicate *part_step* references a predicate called *swap_list*. This has the general form

$$\text{swap_list}(L, i, j)$$

and exchanges the values of L_i and L_j , leaving the remaining elements of the list *L* unchanged. Here is one way to express this in ITL:

$$\forall k < |L|: \\ \left[\text{if } k = i \text{ then } (L_k \leftarrow L_j) \text{ else if } k = j \text{ then } (L_k \leftarrow L_i) \text{ else } (L_k \leftarrow L_k) \right].$$

7.2.1 Correctness of *partition_list*

The correctness of this algorithm can be expressed using two properties. The first one states that the list variable *L*'s final value is a permutation of its initial one:

$$\models \text{partition_list}(L, \text{key}, \text{left_len}) \supset [\text{list_to_bag}(L) \leftarrow \text{list_to_bag}(L)].$$

Here the function $list_to_bag(L)$ equals the bag (multi-set) corresponding to L . Thus we express the fact that the initial and final multi-sets of L are identical.

The second property states that L 's final value is partitioned according to key :

$$\begin{aligned} \models & \text{partition_list}(L, key, left_len) \\ & \supset \text{fin}(\forall i < |L| : [(i < left_len) \equiv (L_i < key)]). \end{aligned}$$

The definition of $partition_list$ uses the predicate $part_loop(L, key, I, J)$ which iterates over L while the variables I and J index to the start and end of the sublist of elements not yet processed. The following property states that $part_loop$ leaves the multi-set representation of L stable:

$$\models \text{part_loop}(L, key, I, J) \supset \text{stable } list_to_bag(L).$$

In addition, throughout the computation the elements L_0, \dots, L_{I-1} are all less than the key and the elements $L_J, L_{J+1}, \dots, L_{|L|-1}$ are greater than or equal to the key:

$$\begin{aligned} \models & \text{part_loop}(L, key, I, J) \supset \\ & [\Box \forall k < I : (L_k < key) \\ & \wedge \Box \forall J \leq k < |L| : (L_k \geq key)]. \end{aligned}$$

Eventual termination of $part_loop$ is indicated by the fact that in every step the value of J is greater than or equal to that of I and either I increases or J decreases by 1. Therefore, the difference $J - I$ continues to decrease by 1:

$$\models \text{part_loop}(L, key, I, J) \supset [(J - I) \text{ gets } (J - I - 1)].$$

Since I is initially 0 and J is initially $|L|$, it follows that $|L|$ units of time are required for the entire computation:

$$\models \text{part_loop}(L, key, I, J) \supset \text{len}(|L|).$$

7.3 Quicksort

Let $sort(e)$ be a function equalling the list expression e 's sorted value. We can then specify the in-place sort of a list variable L by means of the formula

$$L \leftarrow sort(L).$$

Another way to express this is as follows:

$$[list_to_bag(L) \leftarrow list_to_bag(L)] \wedge [fin\ sorted(L)].$$

Here the predicate *sorted*(*e*) is true if the list expression *e* is sorted. Thus, the overall formula states that *L*'s final value is a permutation of its original one since the bag corresponding to *L* remains unchanged. Furthermore, *L*'s final value is in sorted order. One way to sort a list is by using the predicate *serial_quicksort* shown in figure 12. We use a special subscripting construct of the form $e_1[e_2..e_3]$. This is a sublist of the expression e_1 of length $e_3 - e_2$ and has the form

$$\langle e_1[e_2], e_1[e_2 + 1], \dots, e_1[e_3 - 1] \rangle.$$

It can be expressed as follows

$$e_1[e_2..e_3] =_{\text{def}} \langle e_1[i]: e_2 \leq i < e_3 \rangle,$$

where the static variable *i* does not occur freely in e_1 . Note that $e_1[e_3]$ is not included among the sublist's elements. We generally write a sublist expression such as $L[0..k]$ in the form $L_{0..k}$.

The quicksort algorithm given here leaves the list unchanged if it has 0 or 1 elements and otherwise partitions it into two main parts with a pivot element in between. The left half has elements less than the pivot and the right half has elements greater than or equal to the pivot. The left part is then recursively sorted, after which the right part is itself sorted. Throughout this time, the pivot is kept stable. The partitioning operation is performed by the predicate *quick_partition* and the subsequent sorting of the two parts is achieved by *serial_sort_parts*. Both of these predicates are described below.

Here is a program for testing the quicksort predicate:

$$\begin{aligned} &fixed_list(L, 7) \wedge fixed_list(T, 7) \wedge \square display(L, T) \\ &\wedge (L = \langle 4, 5, 2, 0, 6, 1, 3 \rangle) \wedge serial_quicksort(L) \\ &\wedge \forall i < |L|: [T_i \approx (if\ i = L_i\ then\ 1\ else\ 0)]. \end{aligned}$$

The statement *fixed_list*(*L*, 7) specifies that *L* is always a list of length 7. It is logically equivalent to the formula

$$\square list(L, 7)$$

but turns out to be more natural and much more efficient to use in Tempura. We include a list variable *T* that has the same length as *L* and shows which elements

```

define serial_quicksort(L) ≡
  if |L| ≤ 1
  then empty
  else ∃pivot: (
    quick_partition(L, pivot);
    sort_parts(L, pivot)
  ).

define sort_parts(L, k) ≡
  (
    serial_quicksort(L0..k)
    ∧ stable Lk..|L|
  );
  (
    serial_quicksort(L(k+1)..|L|)
    ∧ stable L0..(k+1)
  ).

define quick_partition(L, pivot) ≡
  (
    partition_list(L0..|L|-1, L|L|-1, pivot)
    ∧ stable L|L|-1
  );
  (
    skip ∧ swap_list(L, pivot, |L| - 1)
  ).

```

Figure 12: Predicates for serial quicksort

| | | |
|--------------------------------|-------------------|-------------------|
| State 0: | L=<4,5,2,0,6,1,3> | T=<0,0,1,0,0,0,0> |
| State 1: | L=<1,5,2,0,6,4,3> | T=<0,0,1,0,0,0,0> |
| State 2: | L=<1,5,2,0,6,4,3> | T=<0,0,1,0,0,0,0> |
| State 3: | L=<1,6,2,0,5,4,3> | T=<0,0,1,0,0,0,0> |
| State 4: | L=<1,0,2,6,5,4,3> | T=<0,0,1,0,0,0,0> |
| State 5: | L=<1,0,2,6,5,4,3> | T=<0,0,1,0,0,0,0> |
| State 6: | L=<1,0,2,6,5,4,3> | T=<0,0,1,0,0,0,0> |
| State 7: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 8: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 9: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 10: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 11: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 12: | L=<0,1,2,3,5,4,6> | T=<1,1,1,1,0,0,1> |
| State 13: | L=<0,1,2,3,5,4,6> | T=<1,1,1,1,0,0,1> |
| State 14: | L=<0,1,2,3,5,4,6> | T=<1,1,1,1,0,0,1> |
| State 15: | L=<0,1,2,3,5,4,6> | T=<1,1,1,1,0,0,1> |
| State 16: | L=<0,1,2,3,5,4,6> | T=<1,1,1,1,0,0,1> |
| State 17: | L=<0,1,2,3,4,5,6> | T=<1,1,1,1,1,1,1> |
| Done! Computation length = 17. | | |

Figure 13: Execution of serial quicksort

of L are in their proper positions. For any $i < |L|$, the i -th element of T has the value 1 in a state iff the value of L_i in that state equals i . Of course, this technique only works if L 's initial value is a permutation of the integers $0, 1, \dots, |L| - 1$. Figure 13 shows the program's execution. The list L itself is partitioned from state 0 to state 7 with the value 3 used as the key. In particular, from state 6 to state 7, the value 3 is moved to its proper element L_3 . From state 7 to state 12, the left sublist $L_{0..3}$ (i.e., L_0, L_1 and L_2) is sorted. From state 12 to state 17, the right sublist $L_{4..7}$ (i.e., L_4, L_5 and L_6) is sorted.

7.3.1 Explanation of *quick_partition*

The predicate *quick_partition*($L, pivot$) permutes L and assigns the static parameter *pivot* an index into L so that the elements $L_0, \dots, L_{pivot-1}$ are all less than L_{pivot} and the elements $L_{pivot+1}, \dots, L_{|L|-1}$ are all greater than or equal to L_{pivot} . This is achieved by first invoking the predicate *partition_list* described previously

in section 7.2 on the sublist $L_0, \dots, L_{|L|-2}$ with the rightmost element $L_{|L|-1}$ acting as key. The value of $L_{|L|-1}$ is itself keep stable. Afterwards, the value of $pivot$ is an index to the start of the second half of the partition. The element $L_{|L|-1}$ is then exchanged with L_{pivot} .

7.3.2 Explanation of *serial_sort_parts*

The predicate *serial_sort_parts*($L, pivot$) first sorts the left partition $L_0, \dots, L_{pivot-1}$ and after this the right partition $L_{pivot+1}, \dots, L_{|L|-1}$. During the sorting of either side, the other remains stable. Throughout the entire period the value of the pivot element L_{pivot} is left unchanged since it is already in its proper position.

7.3.3 Parallel quicksort

The predicate *par_quicksort* shown in figure 14 is a parallel variant of the serial algorithm just described. As the execution in figure 15 demonstrates, this version of quicksort can take fewer computational steps than the serial one. The basic change to the sorting technique is seen in the predicate *par_sort_parts*. This sorts the two sublists of L in parallel rather than in succession.

Each sublist has a flag (i.e., *Ready1* or *Ready2*) associated with an invocation of the predicate *sort_process*. This predicate recursively sorts the sublist and then sets the flag to true indicating completion. It then waits for the flag *Done*, which becomes true when both sublists are finished. In the mean time, the sublist is kept stable.

7.4 A Multiplication Circuit

Figure 16 depicts a simple multiplication circuit containing a number of component types such as *flipflop*, *zero_test* and *adder*. This was originally designed and verified by Mike Gordon using the LSM behavioral notation [6]. In figure 17 we define the components as predicates in ITL. Gordon used the boolean values *true* and *false* to represent bit signals. We follow the same convention as can be seen in such devices as *zero_test* and *or_gate*. The components *reg* and *flipflop* in the original specification both provide a form of unit delay and can be modelled by the ITL construct *gets*. The overall multiplier is then given in figure 18.

```

define par_quicksort(L) ≡
  if |L| ≤ 1 then empty
  else ∃pivot: (
    quick_partition(L, pivot);
    par_sort_parts(L, pivot)
  ).

define par_sort_parts(L, pivot) ≡
  ∃Done, Ready1, Ready2: (
    (Done ≈ (Ready1 ∧ Ready2))
    ∧ halt Done
    ∧ quicksort_process(Done, Ready1, L0..pivot)
    ∧ quicksort_process(Done, Ready2, Lpivot+1..|L|)
  )
  ∧ stable Lpivot.

define quicksort_process(Done, Ready, L) ≡
  process(
    [par_quicksort(L) ∧ (Ready ≈ empty)];
    [halt Done ∧ (stable L) ∧ (stable Ready)]
  ).

```

Figure 14: Predicates for parallel quicksort

| | | |
|--------------------------------|-------------------|-------------------|
| State 0: | L=<4,5,2,0,6,1,3> | T=<0,0,1,0,0,0,0> |
| State 1: | L=<1,5,2,0,6,4,3> | T=<0,0,1,0,0,0,0> |
| State 2: | L=<1,5,2,0,6,4,3> | T=<0,0,1,0,0,0,0> |
| State 3: | L=<1,6,2,0,5,4,3> | T=<0,0,1,0,0,0,0> |
| State 4: | L=<1,0,2,6,5,4,3> | T=<0,0,1,0,0,0,0> |
| State 5: | L=<1,0,2,6,5,4,3> | T=<0,0,1,0,0,0,0> |
| State 6: | L=<1,0,2,6,5,4,3> | T=<0,0,1,0,0,0,0> |
| State 7: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 8: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 9: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 10: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 11: | L=<1,0,2,3,5,4,6> | T=<0,0,1,1,0,0,1> |
| State 12: | L=<0,1,2,3,4,5,6> | T=<1,1,1,1,1,1,1> |
| Done! Computation length = 12. | | |

Figure 15: Execution of parallel quicksort

Various internal signals such as *B1* and *L3* are hidden by means of existential quantification (\exists). The signals *in1*, *in2*, *Done* and *Out* are left accessible. The device's structure is represented as a conjunction of instances of the individual components. In addition, equalities are included to properly initialize the signals *Done*, *Out* and *L2*.

In order to test the multiplier, we feed the circuit some data and then terminate when the circuit has the answer. The following program performs these tasks and has the multiplier determine the product of 4 and 9:

$$mult_imp(4, 9, Done, Out) \wedge (\bigcirc halt\ Done) \wedge \square display(Done, Out).$$

The execution of this is shown in figure 19.

7.5 Pulse Generation

We now demonstrate how Tempura can be used to generate, manipulate and display simple digital waveforms. Let us first extend the boolean operators \neg , \wedge and \vee to permit the bit values 0 and 1. For example the value of the expression $\neg 0$ is 1 and the value of $1 \wedge 0$ is 0. We use these constructs in the program in figure 20. This controls the behavior of the four bit variables *W*, *X*, *Y* and *Z*.




Figure 16: Block diagram of multiplication circuit

```

define mux(Switch, In1, In2, Out) ≡
  □ (if Switch then Out = In1 else Out = In2)

define reg(In, Out) ≡ (Out gets In).

define flipflop(In, Out) ≡ (Out gets In).

define dec(In, Out) ≡
  Out ≈ (if In = 0 then 0 else In - 1).

define adder(In1, In2, Out) ≡ (Out ≈ [In1 + In2]).

define zero_test(In, Out) ≡ (Out ≈ [In = 0]).

define or_gate(In1, In2, Out) ≡ (Out ≈ [In1 ∨ In2]).

define zero(Out) ≡ (Out ≈ 0).

```

Figure 17: Definitions of individual components

The signal W is initialized to 0 and then successively oscillates 5 times over 4-unit intervals by means of the following form of sequential iteration:

```

for 5 times do (
  [len(3) ∧ stable W]; [skip ∧ (W ← ¬W)]
).

```

The total length of the period is therefore 20 units. In parallel with this, the signal X is initialized to 0 and then receives the values of W but with unit delay. The same happens from X to Y . The value of Z is always the bit-and of W , X and Y . In figure 21, we display the behavior of the combined system in the form of a timing diagram. This is best viewed when turned sideways. The style of output used here is not hard to generate in Tempura although we omit the details.

7.6 Testing a Latch

The device shown in figure 22 is a simple latch built out of two cross-coupled nor-gates. When the bit inputs S and R are held stable long enough, the outputs Q and \bar{Q} respond to them according to the table in figure 23. If S and R are both

```

define mult_imp(in1, in2, Done, Out) ≡
  ∃ B1, B2, B3, B4, L1, L2, L3, L4, L5, L6, L7, L8, L9, L10 : (
    mux(Done, L9, L8, L7)
    ∧ reg(L7, Out)
    ∧ adder(L9, Out, L8)
    ∧ dec(in1, L6)
    ∧ mux(Done, L6, L4, L5)
    ∧ mux(Done, in1, L3, L1)
    ∧ reg(L1, L2)
    ∧ dec(L2, L3)
    ∧ dec(L3, L4)
    ∧ zero(L10)
    ∧ mux(B4, L10, in2, L9)
    ∧ zero_test(in1, B4)
    ∧ zero_test(L5, B1)
    ∧ zero_test(in2, B2)
    ∧ or_gate(B1, B2, B3)
    ∧ flipflop(B3, Done)
    ∧ (Done = true) ∧ (Out = 0) ∧ (L2 = 0)
  ).

```

Figure 18: Definition of multiplier

| | | | |
|-------------------------------|----|------------|--------|
| State | 0: | Done=true | Out= 0 |
| State | 1: | Done=false | Out= 9 |
| State | 2: | Done=false | Out=18 |
| State | 3: | Done=false | Out=27 |
| State | 4: | Done=true | Out=36 |
| Done! Computation length = 4. | | | |

Figure 19: Simulation of multiplier

```

(W = 0) ∧ (X = 0) ∧ (Y = 0) ∧
for 5 times do (
  [len(3) ∧ (stable W)];
  [skip ∧ (W ← ¬W)]
)
∧ (X gets W) ∧ (Y gets X)
∧ (Z ≈ [W ∧ X ∧ Y]).

```

Figure 20: Program to generate some digital waveforms

| State | 0: | Z | Y | X | W |
|-----------|----|---|---|---|---|
| State 0: | | | | | |
| State 1: | | | | | |
| State 2: | | | | | |
| State 3: | | | | | |
| State 4: | | | | | |
| State 5: | | | | | |
| State 6: | | | | | |
| State 7: | | | | | |
| State 8: | | | | | |
| State 9: | | | | | |
| State 10: | | | | | |
| State 11: | | | | | |
| State 12: | | | | | |
| State 13: | | | | | |
| State 14: | | | | | |
| State 15: | | | | | |
| State 16: | | | | | |
| State 17: | | | | | |
| State 18: | | | | | |
| State 19: | | | | | |
| State 20: | | | | | |

Done! Computation length = 20.

Figure 21: Execution of waveform generator




Figure 22: Block diagram of SR-latch

| operation | S | R | Q | \bar{Q} |
|------------|-----|-----|---------|---------------|
| set to 1 | 1 | 0 | 1 | 0 |
| clear to 0 | 0 | 1 | 0 | 1 |
| no change | 0 | 0 | old Q | old \bar{Q} |
| undefined | 1 | 1 | - | - |

Figure 23: SR-latch operation

```

( $S = 0$ )  $\wedge$  ( $R = 0$ )  $\wedge$  ( $Q = 0$ )  $\wedge$  ( $Qbar = 0$ )  $\wedge$ 
for  $l \in \langle\langle 1, 0 \rangle, \langle 0, 0 \rangle, \langle 0, 1 \rangle, \langle 1, 0 \rangle, \langle 0, 0 \rangle\rangle$ 
do (
   $len(5) \wedge (S \text{ gets } l_0) \wedge (R \text{ gets } l_1)$ 
)
 $\wedge (Q \text{ gets } \neg[R \vee Qbar])$ 
 $\wedge (Qbar \text{ gets } \neg[S \vee Q])$ .

```

Figure 24: Program to simulate SR-latch

0, the device retains its current state. The values of the inputs should not be held simultaneously at 1 since this results in the latch's outputs starting to oscillate.

The program in figure 24 simulates the latch with various values for the inputs S and R . We model each nor-gate as having unit delay. Note that the variable \overline{Q} is referred to as $Qbar$ in the program. The resulting system behavior is displayed in figure 25. The iterative construct

for $v \in e$ do w

is used to sequentially assign a locally scoped variable v the elements of the given list expression e and execute the statement w with each such binding.

7.7 Synchronized Communication

Tempura can be used to model parallel processes that periodically send or receive data from one another. Figure 26 shows a block diagram containing two modules with a line between them. The left module transmits information to the right one by means of a *stream* package that we have implemented in Tempura. We omit the implementation details but give a sample execution of the system in figure 27. In states 3, 6, 9, 12 and 15 a value is passed on. The end of stream marker is passed in state 18. In the final state, the receiver shows that it has successfully accepted the five data messages.

In figure 28 is a block diagram of a system that does parallel lexical analysis, parsing and evaluation of a string of characters representing an arithmetic expression. The heart of the corresponding Tempura program is shown in figure 29.

Here four processes are connected together by means of streams. The first one

| State | 0: | Qbar | Q | R | S |
|-----------|----|------|---|---|---|
| State 0: | | | | | |
| State 1: | | | | | |
| State 2: | | | | | |
| State 3: | | | | | |
| State 4: | | | | | |
| State 5: | | | | | |
| State 6: | | | | | |
| State 7: | | | | | |
| State 8: | | | | | |
| State 9: | | | | | |
| State 10: | | | | | |
| State 11: | | | | | |
| State 12: | | | | | |
| State 13: | | | | | |
| State 14: | | | | | |
| State 15: | | | | | |
| State 16: | | | | | |
| State 17: | | | | | |
| State 18: | | | | | |
| State 19: | | | | | |
| State 20: | | | | | |
| State 21: | | | | | |
| State 22: | | | | | |
| State 23: | | | | | |
| State 24: | | | | | |
| State 25: | | | | | |

Done! Computation length = 25.

Figure 25: Simulation of SR-latch




Figure 26: Block diagram of sender-receiver system

takes the string and feeds it character-by-character into the lexical analyzer. This simultaneously outputs tokens to the parser. This in turn outputs reverse Polish to the evaluator. Eventually the evaluator determines the expression's numeric value. An execution of the system processing the string $10_{\square}+2_{\square}$ is shown in figure 30. We use the character " \square " to visibly represent a space. Note that in states 13 and 24 two communications occur at once.

8 Design of an Interpreter for Tempura

Let us now consider how to built an interpreter for Tempura. This material can be skipped by readers not interested in the intricacies of the system. The aim of the presentation is to give a general idea of how to execute Tempura programs. Therefore a number of low-level aspects of the implementation are not described. Before discussing the design of the system, we look how the approach is applied to the following program:

$$(\bigcirc \bigcirc \bigcirc \text{ empty}) \wedge (I = 0) \wedge (I \text{ gets } I + 1) \wedge \square (J = 2 \cdot I). \quad (6)$$

This is true on intervals of length 4 in which I assumes the successive values 0, 1, 2, 3 and 4 while J simultaneously assumes the values 0, 2, 4, 6 and 8.

One way to execute such a formula is to transform it to a logically equivalent conjunction of two formulas *present_state* and \odot *what_remains*:

$$\textit{present_state} \wedge \odot \textit{what_remains}.$$

Here, the formula *present_state* consists of assignments to the program variables and also indicates whether or not the interval is finished. The formula *what_remains*

```
State 0: Stream_status
State 0: <>
State 1: <>
State 2: <>
State 3: <<0>>
State 4: <>
State 5: <>
State 6: <<10>>
State 7: <>
State 8: <>
State 9: <<20>>
State 10: <>
State 11: <>
State 12: <<30>>
State 13: <>
State 14: <>
State 15: <<40>>
State 16: <>
State 17: <>
State 18: <<>>
State 19: <>
State 20: <>
State 21: <>
State 22: <>
State 23: <>
State 24: <>
State 25: <>
State 26: <>
State 26: output list = <0,10,20,30,40>

Done! Computation length = 26.
```

Figure 27: Execution of sender-receiver system



Figure 28: Block diagram of parallel lexer-parser-evaluator

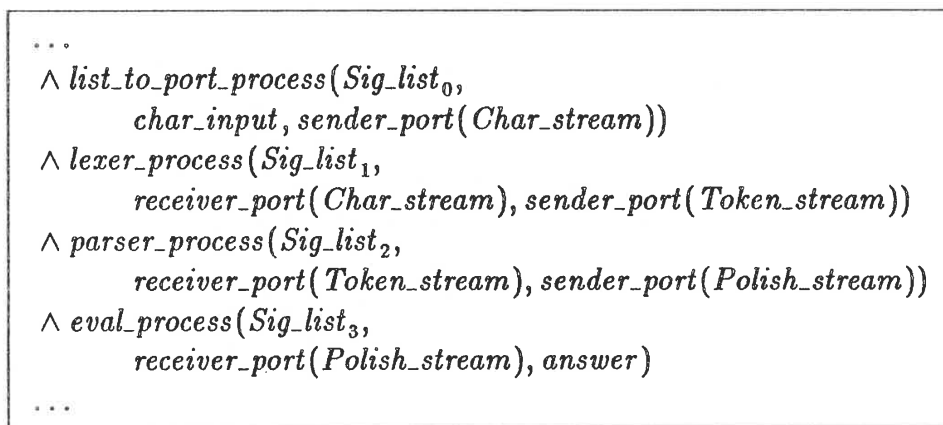


Figure 29: Heart of parallel lexer-parser-evaluator

| State | Char | Token | Polish |
|--------------------------------|-------------------------|---------|---------|
| State 0: | Input string = "10_+2_" | | |
| State 0: | <> | <> | <> |
| State 1: | <> | <> | <> |
| State 2: | <> | <> | <> |
| State 3: | <<"1">> | <> | <> |
| State 4: | <> | <> | <> |
| State 5: | <> | <> | <> |
| State 6: | <<"0">> | <> | <> |
| State 7: | <> | <> | <> |
| State 8: | <> | <> | <> |
| State 9: | <<"_">> | <> | <> |
| State 10: | <> | <> | <> |
| State 11: | <> | <<10>> | <> |
| State 12: | <> | <> | <> |
| State 13: | <<"+">> | <> | <<10>> |
| State 14: | <> | <> | <> |
| State 15: | <> | <<"+">> | <> |
| State 16: | <> | <> | <> |
| State 17: | <<"2">> | <> | <> |
| State 18: | <> | <> | <> |
| State 19: | <> | <> | <> |
| State 20: | <<"_">> | <> | <> |
| State 21: | <> | <> | <> |
| State 22: | <> | <<2>> | <> |
| State 23: | <> | <> | <> |
| State 24: | <<>> | <> | <<2>> |
| State 25: | <> | <> | <> |
| State 26: | <> | <<>> | <> |
| State 27: | <> | <> | <> |
| State 28: | <> | <> | <<"+">> |
| State 29: | <> | <> | <> |
| State 30: | <> | <> | <> |
| State 31: | <> | <> | <<>> |
| State 32: | <> | <> | <> |
| State 33: | <> | <> | <> |
| State 34: | <> | <> | <> |
| State 34: | Final answer = 12 | | |
| Done! Computation length = 34. | | | |

Figure 30: Execution of lexer-parser-evaluator system

| |
|--|
| <p>Before state 0: $(\circ\circ\circ \text{ empty}) \wedge (I = 0) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)$</p> <p>After state 0: $[(I = 0) \wedge (J = 0) \wedge \text{ more}]$ $\wedge \odot[(\circ\circ \text{ empty}) \wedge (I = 1) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$</p> <p>Before state 1: $(\circ\circ \text{ empty}) \wedge (I = 1) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)$</p> <p>After state 1: $[(I = 1) \wedge (J = 2) \wedge \text{ more}]$ $\wedge \odot[(\circ \text{ empty}) \wedge (I = 2) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$</p> <p>Before state 2: $(\circ \text{ empty}) \wedge (I = 2) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)$</p> <p>After state 2: $[(I = 2) \wedge (J = 4) \wedge \text{ more}]$ $\wedge \odot[\text{empty} \wedge (I = 3) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$</p> <p>Before state 3: $\text{empty} \wedge (I = 3) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)$</p> <p>After state 3: $[(I = 3) \wedge (J = 6) \wedge \text{ empty}]$ $\wedge \odot[\text{false} \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$</p> |
|--|

Figure 31: Transformation of formula (6)

is what is executed in subsequent states if the interval does indeed continue on. Thus, it can be viewed as a kind of continuation.

For the formula under consideration, the value of *present_state* is as follows:

$$(I = 0) \wedge (J = 0) \wedge \text{ more.}$$

The value of *what_remains* is the formula

$$(\circ\circ \text{ empty}) \wedge (I = 1) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I).$$

In figure 31 we show the effects of such transformations before and after each of the four states of the execution.

8.1 Variables Used by the Interpreter

The operation of the Tempura interpreter is based on the technique just described. When executing a Tempura program, the interpreter makes use of the four variables *Program*, *Memory*, *Current_Env* and *Current_Done_Cell*. They are accessed by the interpreter's various routines. Below is a summary of each of these variables:

- *Program*: This is a variable that initially contains the Tempura program itself. After the execution of each state, it is transformed to a formula of the form $\odot w$, where w describes what should be done in the next state.
- *Memory*: This is an indexed list of cells. Each cell can be empty or contain a value such as an integer or a list descriptor. At the beginning of each state, every cell is assigned the empty list $\langle \rangle$, thus indicating that no value is being stored. When a value c is to be placed in the cell, the cell's actual contents are assigned the singleton $\langle c \rangle$.
- *Current_Env*: This contains an environment, which is a list having a separate entry for each variable in the Tempura program. Each entry is a pair with the name of the associated variable as well as an index to a memory cell that holds the variable's value. For the Tempura program under consideration, the value of *Current_Env* might be the following:

$$\langle \langle "I", 0 \rangle, \langle "J", 1 \rangle \rangle.$$

- *Current_Done_Cell*: This equals the index of a memory cell called the *done-flag*. The Tempura program places either *true* or *false* in the done-flag during every state, thus indicating whether or not that state is the final one. For example, in a state where the statement *empty* is encountered, the interpreter puts the value *true* in the done-flag cell. If the statement *more* is encountered, the value *false* is used instead. If the user fails to assign a value to the done-flag in a particular state, the interpreter detects this and flags an error.

8.2 Basic Execution Algorithm

The basic algorithm used by the interpreter can be represented in the following procedural form:

```
begin
  local Program, Memory, Current_Env, Current_Done_Cell;
  prepare_execution_of_program;
  loop
    execute_single_state
    exit when Memory[Current_Done_Cell] = (true)
    otherwise
      advance_to_next_state
  end.
```

Here is a more detailed look at each part of the process:

- *prepare_execution_of_program*: The interpreter's variable *Program* is assigned the program's syntax tree and the variable *Current_Env* is initialized to indicate suitable references into the memory. The memory itself is allocated to have one cell for each program variable as well as a cell whose index is placed in *Current_Done_Cell*. All the memory cells are initially emptied (i.e., set to $\{\}$).
- *execute_single_state*: The value of the variable *Program* is transformed until it is of the form $\odot w$. A check is made to ensure that the done-flag indexed by *Current_Done_Cell* has been set to *true* or *false*. The actual transformations used for each Tempura construct are described later on. All assignments occurring in the current state are reflected in the values of the memory's cells.
- *advance_to_next_state*: If the current state is not the last, preparations are made for processing the next one. This is done by emptying the contents of all the memory cells and deleting the leading operator \odot from the formula held in *Program*.

Figure 32 shows the details of executing the simple program (6).

Current_Env: $\langle\langle "I", 0 \rangle, \langle "J", 1 \rangle\rangle$,
Current_Done_Cell: 2.

Before state 0:
Program: $[(\bigcirc \bigcirc \bigcirc \text{empty}) \wedge (I = 0) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$,
Memory: $\langle\langle \rangle, \langle \rangle, \langle \rangle\rangle$.

After state 0:
Program: $\odot[(\bigcirc \bigcirc \text{empty}) \wedge (I = 1) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$,
Memory: $\langle\langle 0 \rangle, \langle 0 \rangle, \langle \text{false} \rangle\rangle$.

Before state 1:
Program: $[(\bigcirc \bigcirc \text{empty}) \wedge (I = 1) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$,
Memory: $\langle\langle \rangle, \langle \rangle, \langle \rangle\rangle$.

After state 1:
Program: $\odot[(\bigcirc \text{empty}) \wedge (I = 2) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$,
Memory: $\langle\langle 1 \rangle, \langle 2 \rangle, \langle \text{false} \rangle\rangle$.

Before state 2:
Program: $[(\bigcirc \text{empty}) \wedge (I = 2) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$,
Memory: $\langle\langle \rangle, \langle \rangle, \langle \rangle\rangle$.

After state 2:
Program: $\odot[\text{empty} \wedge (I = 3) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$,
Memory: $\langle\langle 2 \rangle, \langle 4 \rangle, \langle \text{false} \rangle\rangle$.

Before state 3:
Program: $[\text{empty} \wedge (I = 3) \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$,
Memory: $\langle\langle \rangle, \langle \rangle, \langle \rangle, \langle \rangle\rangle$.

After state 3:
Program: $\odot[\text{false} \wedge (I \text{ gets } I + 1) \wedge \square(J = 2 \cdot I)]$,
Memory: $\langle\langle 3 \rangle, \langle 6 \rangle, \langle \text{true} \rangle\rangle$.

Figure 32: Details of execution of formula (6)

8.3 Description of the Procedure *execute_single_state*

Here is now the definition of the procedure *execute_single_state*:

```
procedure execute_single_state:
begin
  while  $\neg is\_reduced\_stmt(Program)$  do
    transform_stmt(Program);
  if Memory[Current_Done_Cell] =  $\langle \rangle$  then
    error("Termination status not specified.")
end.
```

The while-loop repeatedly transforms the program until the test

$$is_reduced_stmt(Program)$$

is true, signifying that the program is reduced to the form $\odot w$. Each iteration of the loop corresponds to one *pass* over the program. Sometimes many passes can be required to completely process one state. Once this is achieved, no further transformations take place in the current state. As long as the program is not in reduced form, the procedure *transform_stmt* is invoked to scan it. The operation of *transform_stmt* is explained in detail below. Once the program is fully reduced, a check is made to ensure that the done-flag has been properly set.

8.4 Description of the Procedure *transform_stmt*

The procedure *transform_stmt* has the form

$$transform_stmt(Statement).$$

It makes a single left-to-right pass over a Tempura statement held in the parameter *Statement* and transforms it, while simultaneously extracting information about Tempura variables and termination. The contents of the memory cells are updated in the process. Furthermore, when a Tempura *request* or *display* statement is reduced, communication with the user takes place.

As noted in the presentation of Tempura's syntax, a formula such as

$$(I = 2) \wedge (J = I + 3)$$

can be viewed as both a statement and as a boolean test. In addition, a variable such as I can be considered either a location or an expression. Therefore, there exist two other reduction routines. The procedure *transform_loc* is for locations and the procedure *transform_expr* is for expressions. These are both used by *transform_stmt* and are described in sections 8.5 and 8.6, respectively.

Let us now consider the behavior of *transform_stmt* on individual types of Tempura statements.

8.4.1 Implementing the statements *true* and *false*

When the statement *true* is encountered, it is immediately reduced to the form $\odot true$. The values of the memory cells are not affected. Thus, *true* can be thought of as a no-operation statement. On the other hand, when *false* is encountered as a statement, the interpreter terminates execution with an error. This provides the user a way to abnormally stop the program if inconsistencies are detected.

8.4.2 Implementing equalities

An equality has the general form $l = e$, where l is a location and e is an expression. This is executed by first transforming l and e to new forms l' and e' . See sections 8.5 and 8.6 for details. If either l' or e' is not yet fully reduced, the statement

$$l' = e'$$

is returned as the result. If both are successfully reduced, then e' is a constant c and l' has one of the following two forms:

$$@loc(k), \quad \circ l''.$$

The construct $@loc(k)$ is an internal location descriptor that references the memory cell indexed by the integer k . If l' is such a descriptor, the equality is processed by placing the constant c in the k -th cell. A check is made to ensure that the cell is empty prior to the assignment. The overall statement is then reduced to the form $\odot true$. This is returned as the result.

If the location l' is of the form $\circ l''$ then the assignment is transformed to the conditional statement

$$\text{if more then } \odot(l'' = c) \text{ else false.}$$

This postpones the actual assignment to the next state. The *more* construct is used to ensure that the interval does indeed continue. If the test fails, the statement *false* is executed, thus resulting in an error.

8.4.3 Implementing the statements *empty* and *more*

The transformation of the statement *empty* places the value *true* in the done-flag currently referenced by *Current_Done_Cell*. The statement is then reduced to $\textcircled{\text{w}}\textit{false}$. The transformation of *more* places the value *false* in the done-flag referenced by *Current_Done_Cell* and then reduces to $\textcircled{\text{w}}\textit{true}$. In the case of both statements, a check is made that the done-flag is empty (i.e., equals $\langle \rangle$) prior to the assignment.

8.4.4 Implementing the *request* and *display* statements

A statement of the form

$$\textit{request}(l_1, \dots, l_n)$$

is processed by first prompting the user for n values c_1, \dots, c_n and then transforming the statement to a conjunction of equalities

$$(l_1 = c_1) \wedge \dots \wedge (l_n = c_n).$$

This is then re-reduced.

During the execution of a statement

$$\textit{display}(e_1, \dots, e_n),$$

the expressions e_1, \dots, e_n are all transformed. If any fail to completely reduce, the statement is returned unchanged. Otherwise, the resulting values are displayed and the statement is reduced to $\textcircled{\text{w}}\textit{true}$.

8.4.5 Implementing conjunctions

A conjunction $w_1 \wedge w_2$ is processed by first reducing both operands to forms w'_1 and w'_2 . If this is successful, the statement w'_1 will be of the form $\textcircled{\text{w}}w''_1$ and the statement w'_2 will be of the form $\textcircled{\text{w}}w''_2$. The overall statement can then be reduced

to the form $\textcircled{w}(w_1'' \wedge w_2'')$. If either w_1'' or w_2'' is *true*, it can be omitted from the result. If w_1' or w_2' is not yet fully reduced, the conjunction $w_1' \wedge w_2'$ is returned instead.

8.4.6 Transforming the statements \textcircled{w} , $\bigcirc w$ and $\square w$

A statement of the form $\textcircled{w} w$ is already in reduced form and requires no further processing. Statements of the forms $\bigcirc w$ and $\square w$ are rewritten using the following equivalences and then reprocessed:

$$\begin{aligned}\bigcirc w &\equiv \text{more} \wedge \textcircled{w} w \\ \square w &\equiv w \wedge \textcircled{w} \square w.\end{aligned}$$

8.4.7 Implementing implications

A conditional statement of the form $b \supset w$ is treated by first reducing the boolean expression b to either *true* or *false*. If b reduces to *true*, the overall statement is changed to w and again transformed. If b reduces to *false*, the result is $\textcircled{w} \text{true}$.

8.4.8 Implementing some other statements

As was mentioned earlier, statements using the operators \bigcirc , *if*, *gets*, *stable* and *halt* can be converted to equivalent forms using the operators described above. For example, a statement $l \text{ gets } e$ is first expanded to the form

$$\square [\text{more} \supset ((\bigcirc l) = e)]$$

and then re-reduced. Execution efficiency can be greatly increased if each new operator has its own specialized transformation sequence.

8.5 Implementing Locations

A location that is a variable such as I or J is transformed by converting it to an internal descriptor $\text{@loc}(k)$, where k is the variable's cell number as indexed by the environment contained in *Current_Env*. Locations can also be of the form $\bigcirc l$, where l is itself a location. These are considered to already be in reduced form.

8.6 Implementing Expressions

An expression is in reduced form if it is either an arithmetic or boolean constant. Since constants such as 3 and *true* are already in reduced form, they require no further processing. A variable such as *I* is handled by first reducing it as a location. The result is an internal location descriptor of the form $@loc(k)$. The actual value of *k* is determined by looking at *I*'s entry in the current environment. The descriptor is immediately re-reduced as an expression.

A location descriptor $@loc(k)$ is itself transformed by examining the memory cell indexed by the integer *k*. If that cell is empty (i.e., equal to $\langle \rangle$), the descriptor is returned unchanged. Otherwise, the value stored in the cell is returned instead of the descriptor. For example, if the cell equals the singleton $\langle 4 \rangle$, the constant 4 is the result.

An expression such as the sum $e_1 + e_2$ is handled by first reducing both e_1 and e_2 and then adding them. A conditional expression of the form *if b then e₁ else e₂* is reduced in the same way as a conditional statement. Boolean expressions such as $\neg b$ and $b_1 \wedge b_2$ can be expanded to conditional expressions and then re-reduced. The boolean expressions *empty* and *more* are transformed to the constructs $@loc(done-cell)$ and $\neg @loc(done-cell)$, respectively, where *done-cell* is the value of *Current_Done_Cell*. The new expressions are then immediately re-reduced.

8.7 Implementing Static Variables, Lists and Quantifiers

We now consider how to extend the interpreter to handle static variables, lists and quantifiers. The operators *fin* and \leftarrow as well as predicates and functions are also discussed.

8.7.1 Implementing static variables

Static variables are implemented by altering the memory slightly to include an extra boolean flag in every cell. This is set to *true* if the cell has been designated as static. When a memory cell is allocated for a state variable, the cell's flag is set to *false*. While clearing memory, the interpreter only empties those cells that have the flag equalling *false*. Thus once a value is put in a static cell, it is never lost.

8.7.2 Implementing lists

The value of a list is represented as a special descriptor of the form

$$@list(length, offset).$$

Here *length* equals the number of list elements and *offset* is the index of the first of a series of consecutive memory cells for storing the elements.

8.7.3 Implementing the operators *list*, *fixed_list* and *stable_struct*

A statement of the form $list(l, e)$ is processed by first reducing the location l and expression e . A series of consecutive memory cells are then allocated, one for each list element. If the cell referenced by location l is designated static, then so are these cells. Afterwards, l 's cell is assigned a suitable list descriptor.

The implementation of the construct $fixed_list(l, e)$ mentioned in section 7.3 is similar to that of $list(l, e)$. However, the memory cells initially allocated for the list elements continue to be used during the entire interval. This is done by simply reassigning the list description stored in the location l from state to state. Thus a statement such as

$$fixed_list(Table, 20)$$

is much more efficient than the logically equivalent form

$$\square list(Table, 20).$$

In addition, by using $fixed_list$, we ensure that the locations of the elements of *Table* do not change over time. This turns out to be very important when passing them as parameters to temporal predicates.

The construct $stable_struct(l)$ mentioned in section 7.1.1 has an implementation similar to that for $fixed_list$ in that the list descriptor stored in l is repeatedly reassigned throughout execution. However, the list elements of l are assumed to be already allocated. The conjunction

$$list(l, e) \wedge stable_struct(l)$$

can in fact be used to implement the statement

$$fixed_list(l, e).$$

8.7.4 Implementing list constructors

A list constructor of the form

$$(e_0, \dots, e_{n-1})$$

is evaluated by first reducing the expressions e_0, \dots, e_{n-1} and then allocating n consecutive static memory cells in which the values of e_0, \dots, e_{n-1} are stored. A suitable list descriptor is then returned as the overall value of the expression. The iterative list constructor is implemented in a similar way.

8.7.5 Implementing subscripts

A subscripted location $l[e]$ is processed by first evaluating the location l and the expression e . The value of l 's cell is fetched and checked to ensure that it is a list. Furthermore, a check is made to ensure that e 's value is within the list's range. A location of the form $@loc(k)$ is then returned, where k is the sum of the list's offset into memory and e 's value.

An expression $e_1[e_2]$ is treated in a similar manner. However once the relevant element's location has been determined, the value in the corresponding memory cell is fetched.

The sublist construct $l[e_1 .. e_2]$ mentioned in section 7.3 is implemented by first reducing the location l and the expressions e_1 and e_2 . We then form a list descriptor with offset to the element in l referenced by e_1 and with length $e_2 - e_1$. This descriptor is stored in a new cell and the location of that cell is the result of the reduction.

8.7.6 Implementing the list-length operator

An expression $|e|$ is reduced by first determining the list descriptor of the expression e and then returning the length field stored in it.

8.7.7 Implementing existential quantifiers

A statement of the form

$$\exists V_1, \dots, V_n: w$$

is processed by creating a new environment env containing entries for the quantified variables V_1, \dots, V_n . Each of these variables has a fresh memory cell allocated for it. The entries in env for other variables are the same as in the surrounding environment. The statement is then transformed to an interval construct of the form

$$@exists(env, w).$$

This is immediately re-reduced. Such a statement is evaluated by saving the contents of $Current_Env$, setting $Current_Env$ to env and then transforming the statement w . Afterwards, the old contents of $Current_Env$ are restored. If the transformation result w' is successful, it has the form $\odot w''$. Therefore, the overall statement is converted to the form

$$\odot[@exists(env, w'')]$$

and returned. If w' is not fully reduced, the overall statement is given the form

$$@exists(env, w')$$

and returned.

8.7.8 Implementing bounded universal quantification

A statement of the form

$$\forall v < e: w$$

can be readily handled by first reducing the expression e to a constant c and then building a conjunction of the form

$$(\exists v: [(v = 0) \wedge w]) \wedge (\exists v: [(v = 1) \wedge w]) \wedge \dots \wedge (\exists v: [(v = c - 1) \wedge w]).$$

The conjunction is immediately re-reduced.

8.7.9 Implementing the operators len , fn and \leftarrow

The implementation of a statement $len(e)$ first reduces the expression e to another expression e' . If this is not a constant, the overall statement is changed to $len(e')$ and returned. On the other hand, if e' is a constant c , the statement is reduced using the equivalence

$$len(c) \equiv \text{if } (c = 0) \text{ then empty else } \odot len(c - 1).$$

A statement of the form $fin\ w$ is rewritten as the logically equivalent formula

$$if\ empty\ then\ w\ else\ (\oplus\ fin\ w)$$

and then immediately re-reduced. A temporal assignment $l \leftarrow e$ is processed by first reducing the expression e to a constant c . The overall statement is then changed to the form

$$fin(l = c).$$

This is immediately re-reduced.

8.7.10 Implementing predicate and function definitions

A predicate definition has the form

$$define\ p(V_1, \dots, V_n) \equiv \hat{w}.$$

We execute this by creating a special descriptor of the form

$$@predicate(env, \langle V_1, \dots, V_n \rangle, w),$$

where env is the current contents of the environment variable *Current_Env*. The descriptor is stored in the predicate variable p 's memory cell and provides enough information to properly access the predicate when it is invoked. The overall predicate definition is then reduced to the form $\oplus\ true$. Function definitions are similarly handled.

8.7.11 Implementing predicate and function invocations

A predicate invocation has the form

$$p(e_1, \dots, e_n).$$

We first access the value in p 's memory cell and determine that it is indeed a special descriptor of the form

$$@predicate(env, \langle V_1, \dots, V_n \rangle, w).$$

The locations of actual parameters e_1, \dots, e_n are then reduced. If any are in fact constants or expressions, their values are stored in freshly allocated static memory cells. A new environment env' is created in which each formal parameter

V_1, \dots, V_n has an entry pointing to the corresponding actual parameter's location. The entries for other variables are made identical to those found in the predicate-descriptor's environment env .

We now transform the predicate invocation to the statement

$$@call(env', w),$$

where env' is the environment just constructed and w is the statement contained in the predicate-descriptor. The internal operator $@call$ is then processed by saving the current contents of the variable $Current_Env$, setting $Current_Env$ to env' , reducing w and then restoring $Current_Env$'s old value. If the transformed statement w' is completely reduced, then it has the form $\odot w''$. The overall invocation is then rewritten as follows:

$$\odot[@call(env', w'')].$$

On the other hand, if the transformation is not complete, the statement

$$@call(env', w')$$

is returned so that this can be executed during the current state's next pass.

Function invocations are similarly handled. They are eventually reduced to something of the form

$$@call(env, c),$$

where env is an environment and c is a constant. At this time, we discard the operator $@call$ and the environment and transform the overall expression to be simply the constant c .

8.8 Implementing *chop* and Iterative Operators

We now turn to the implementation of *chop* and related operators such as *for* and *while*.

8.8.1 Implementing the operator *chop*

The *chop* construct has the form

$$w_1; w_2,$$

where w_1 and w_2 are formulas. The first step in processing this is to allocate a memory cell to serve as the local done-flag associated with the subinterval in which w_1 is executed. The *chop* statement is then transformed to the internal form

$$@chop(done-cell, w_1, w_2).$$

Here *done-cell* is the index of the memory cell serving as the local done-flag.

The construct *@chop* is executed by first saving the current value of the variable *Current_Done_Cell* and setting *Current_Done_Cell* to the index *done-cell*. The statement w_1 is transformed in this context. Afterwards, we restore the old value of *Current_Done_Cell*. If the transformation's result w' is not yet fully reduced, the *@chop* statement is altered to have the form

$$@chop(done-cell, w', w_2)$$

and then returned as the overall result. However, if w' is fully reduced, it has the form $\odot w''$. The *@chop* statement is therefore transformed to the following conditional form:

$$if @loc(done-cell) then w_2 else \bigcirc @chop(done-cell, w'', w_2).$$

This is immediately re-reduced. Thus, if the value of the local done-flag is true, the subinterval in which w_1 was executed is empty and therefore w_2 can be started right away. If the local done-flag is false, then the execution of the *@chop* statement is continued to the next state of the overall interval. Note that the operator *strong-next* is used to indicate that the overall interval is not yet finished.

8.8.2 Implementing for-loops and other iterative statements

A statement of the form

$$for\ e\ times\ do\ w$$

is transformed by first reducing the expression e to a constant c . The overall statement is then rewritten as the conditional form

$$if\ c = 0\ then\ empty\ else\ (w; [for\ c - 1\ times\ do\ w])$$

and again reduced.

The implementation of an indexed for-loop of the form

$$for\ v < e\ do\ w$$

first reduces the expression e to a constant c . The statement is then transformed to the related construct

$$\text{for } 0 \leq v < c \text{ do } w.$$

This is expanded using the following rule:

$$\begin{aligned} \text{for } c_1 \leq v < c_2 \text{ do } w &\equiv_{\text{def}} \\ &\text{if } c_1 \geq c_2 \text{ then empty} \\ &\text{else } ([\exists v: (v = c_1 \wedge w)]; [\text{for } c_1 + 1 \leq v < c_2 \text{ do } w]). \end{aligned}$$

A while-loop is reduced using the equivalence

$$\text{while } b \text{ do } w \equiv (\text{if } b \text{ then } [w; \text{while } b \text{ do } w] \text{ else empty}).$$

The iterative constructs *repeat* and *loop* are treated similarly.

8.8.3 Implementing the operator *skip*

The operator *skip* is first transformed to the form

$$\bigcirc \text{ empty}$$

and then re-reduced.

8.9 Alternative Interpreters

The interpreter that we have described represents one technique for executing Tempura programs. It is rather easy to understand but suffers from being relatively slow. Let us now consider some alternative approaches and features. Most of them increase execution efficiency at the expense of generality. With a proper mix of these techniques, we feel that we can achieve speeds comparable with conventional imperative programming languages.

8.9.1 Immediate assignments

When the interpreter encounters the equality

$$(\bigcirc I) = 1$$

it postpones the actual assignment to the next state by using the form

$$\oplus(I = 1).$$

An alternate approach involves performing the assignment immediately. This can lead to increased execution efficiency. However, one drawback is that statements with parallel assignments such as the following are not properly handled:

$$([\circ I] = I + 1) \wedge ([\circ J] = J + I).$$

In this example, we would alter I 's value before being able to compute the expression $J + I$. Thus, the next value of J would be unknown. One way to get around this is to reorder the statement as follows:

$$([\circ J] = J + I) \wedge ([\circ I] = I + 1).$$

8.9.2 Two-level memory

So far, the interpreter we have presented maintains a single data value for each location. It is often attractive to maintain two such values: one for the current state and one for the next state. This permits the technique of immediate assignment described above to properly work on parallel assignments such as

$$([\circ I] = I + 1) \wedge ([\circ J] = J + I).$$

Note that static locations still only require a single value.

8.9.3 Single-pass processing

The current interpreter can make a number of left-to-right passes over a statement within each executed state. For example, the compound statement

$$(J = I + 2) \wedge (I = 1)$$

requires two passes. During the first pass the value for I is determined and during the second pass the value for J is determined. On the other hand, the following statement can be completely reduced in a single pass:

$$(I = 1) \wedge (J = I + 2).$$

If we use an interpreter having immediate assignment, then many useful programs can be written which require only one pass per executed state. Once they are debugged, such programs have the potential of being executed much faster since the interpreter need not do the extra processing required to detect and handle multiple passes. For example, when a memory cell is read, no test need be made regarding whether a value is already stored.

8.9.4 Time stamps

A time stamp can be included as part of each memory cell. Whenever the cell is stored into, the current state number is included with the data. When the cell is accessed, its state number can be checked to ensure that the data is current and not left over from a previous state. This approach eliminates the need to empty all cells at the beginning of every state. Many debugged programs require only one pass per executed state and can therefore be executed without the interpreter having to regularly empty memory cells or use time stamps.

8.9.5 Ignoring the operator *stable*

As we have noted, many debugged Tempura programs can be run using one pass per state and without emptying memory cells or maintaining time stamps. Any statement of the form *stable l* in such programs can be ignored since the value of the location *l*'s memory cell will automatically remain unchanged. If *l* is a list, then the effect will safely propagate to its elements. Thus, the savings gained by not processing *stable* can be considerable. This technique should be used with great care since time-dependent errors can go undetected.

8.9.6 Redundant assignments

The interpreter does not permit two or more redundant assignments to the same location to occur in a single state. For example, the statement

$$(I = 1) \wedge (I = 1)$$

causes an error in processing even though it is logically equivalent to the statement

$$I = 1.$$

The principle of nonredundancy also applies to interval termination constructs such as *empty*. Only one such construct should be used in each state of an interval. Thus a formula such as

$$empty \wedge empty,$$

although logically consistent, results in an error when executed. Tempura constructs such as the operators *empty*, *more*, *halt* and \circ affect interval length. Constructs such as \wedge , \oplus and \square do not in themselves specify anything about termination.

If desired, the interpreter can be modified to permit the kind of redundant assignments mentioned here. We currently have reservations but more experience is needed to resolve this issue.

8.9.7 Special-purpose constructs

Various constructs can be added to Tempura in order to speed execution. For instance, the operator $:=$ is called *unit assignment* and is similar to temporal assignment (\leftarrow). However, unit assignment only works in intervals of length 1 and turns out to be more efficiently implementable. A statement $l := e$ is first transformed to the form

$$(\circ l) = e$$

and then re-reduced. In addition, a test is made to ensure that the interval is indeed of unit length.

8.9.8 Suppressing checks

Most programming systems have facilities for checking subscript ranges, detected undefined variables and performing other such tests. The Tempura interpreter includes additional consistency checks regarding temporal behavior. For example, a statement such as

$$(I = 1) \wedge (I = 2)$$

is erroneous since the same variable cannot receive two values in one state. Similarly, the interpreter detects an error in the following program because no value is specified for the variable I in the second state:

$$(I = 1) \wedge [skip; (skip \wedge [I \leftarrow I + 1])] \wedge \square display(I).$$

These types of temporal checks can be suppressed. This increases the speed of execution but of course can result in various time-dependent bugs going unnoticed.

8.9.9 Call-by-name

When a function or predicate is invoked, the interpreter determines the locations of all actual parameters and uses them to build an environment. The effect is to implement parameter passing using call-by-reference. As in other programming languages, this is more efficient than call-by-name but has certain drawbacks. For example, when a subscripted element such as $L[I]$ occurs as a parameter, its location is determined only once. Therefore any change in the value of the index I is not reflected in the corresponding formal parameter. Our experience is that most subscripts tend to be static expressions (e.g., $L[1]$ or $L[i]$) so this is not a great limitation. A more subtle problem occurs if one specifies that L is a list using a statement such as

□ *list(L, 3).*

In every state, new memory cells are allocated for L 's three elements. Therefore, a reference to, say, $L[2]$ will change over time. Thus, if $L[2]$ is passed as a parameter to a temporal predicate, the reference used will become obsolete after the first state. This is one reason we use the constructs *fixed_list* and *stable_struct* for creating and maintaining such lists. Both of them ensure that the locations of list elements do not move around.

If one desires, it is not hard to alter the interpreter to implement parameter passing using call-by-name. This seems to be a trade off between speed and generality. It should be noted that the current interpreter permits one to simulate call-by-name by means of parameterless functions. For example, suppose we define the function f and the predicate p as follows:

define f() = L[I],
define p(g) ≡ (g() gets [g() + 1]).

The predicate invocation $p(f)$ is therefore equivalent to the statement

$L[I]$ gets $(L[I] + 1)$.

9 Experimental Features

In this section we examine some experimental constructs that are not so well understood, yet have interesting applications and properties. The first involves the concept of temporal projection. Following this is a brief look at lambda expressions and their application to representing pointers. Finally, we discuss the *process* construct and its use in specifying parallel programs.

9.1 Temporal Projection

When modelling hardware, it is natural to look at a circuit's behavior at different granularities of time. For example, the units of time might correspond to nanoseconds or to clock ticks. We use the term *temporal projection* to denote the process of mapping from one level of time to another. In earlier work [8,25] we looked at one way to add operators for temporal projection to ITL. Since then we have developed an approach which is a bit easier to use and can be readily incorporated into Tempura. The main new construct is a formula of the form

$$w_1 \text{ proj } w_2,$$

where w_1 and w_2 are themselves formulas. This has the following semantics:

$$\begin{aligned} \mathcal{M}_\sigma \llbracket w_1 \text{ proj } w_2 \rrbracket = \text{true} \quad &\text{iff} \\ &\text{for some relative times } \tau_0, \tau_1, \dots, \tau_m \in \{0, 1, \dots, |\sigma|\}, \text{ the following are true:} \\ &0 = \tau_0 \leq \tau_1 \leq \dots \leq \tau_m = |\sigma|, \\ &\text{and } \mathcal{M}_{\langle \sigma_{\tau_i}, \sigma_{\tau_{i+1}}, \dots, \sigma_{\tau_{i+1}} \rangle} \llbracket w_1 \rrbracket = \text{true}, \text{ for all } i < m, \\ &\text{and } \mathcal{M}_{\langle \sigma_{\tau_0}, \sigma_{\tau_1}, \dots, \sigma_{\tau_m} \rangle} \llbracket w_2 \rrbracket = \text{true}. \end{aligned}$$

This definition makes the formula $w_1 \text{ proj } w_2$ true on any interval meeting two conditions. First, the interval can be broken up into a series of consecutive subintervals each satisfying w_1 . Second, the interval formed from the string of end-states of these subintervals satisfies w_2 .

Consider, for example, the formula

$$\text{len}(2) \text{ proj } [\text{len}(4) \wedge (I = 0) \wedge (I \text{ gets } I + 1)].$$

This is true on any interval σ whose length is 8 and in which I 's value starts at 0 and increases by 1 from σ_0 to σ_2 , from σ_2 to σ_4 and so forth. To show this,

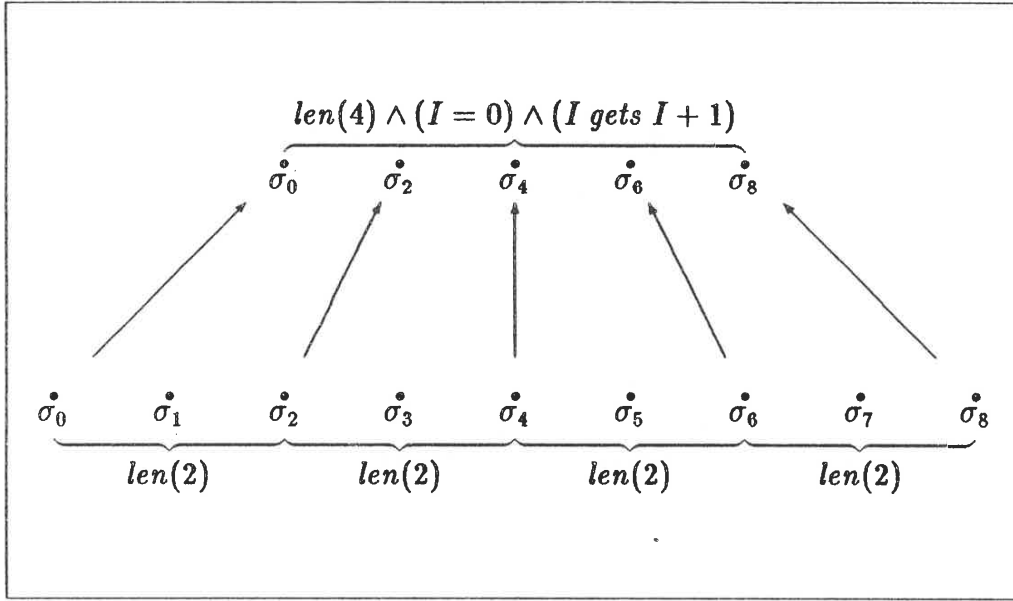


Figure 33: Example of projection

we use the definition of *proj* with m equalling 4 and with τ 's elements having the assignments

$$\tau_0 = 0, \quad \tau_1 = 2, \quad \tau_2 = 4, \quad \tau_3 = 6, \quad \tau_4 = 8.$$

Thus the left operand $len(2)$ is true on each of the following subintervals of σ :

$$\langle \sigma_0 \sigma_1 \sigma_2 \rangle, \quad \langle \sigma_2 \sigma_3 \sigma_4 \rangle, \quad \langle \sigma_4 \sigma_5 \sigma_6 \rangle, \quad \langle \sigma_6 \sigma_7 \sigma_8 \rangle.$$

The right operand

$$len(4) \wedge (I = 0) \wedge (I \text{ gets } I + 1)$$

is true on the projected interval

$$\sigma_0 \sigma_2 \sigma_4 \sigma_6 \sigma_8.$$

Figure 33 illustrates the projection pictorially. The value of I in odd-numbered states is not specified. This example is in fact logically equivalent to the formula

$$(I = 0) \wedge \text{for 4 times do } (len(2) \wedge [I \leftarrow I + 1]).$$

Note that different projections can be combined in parallel as well as with other kinds of formulas.

| |
|--------------------------------|
| State 0: M= 1 |
| State 1: M= 1 |
| State 2: M= 1 |
| State 3: M= 2 |
| State 4: M= 2 |
| State 5: M= 2 |
| State 6: M= 4 |
| State 7: M= 4 |
| State 8: M= 4 |
| State 9: M= 8 |
| State 10: M= 8 |
| State 11: M= 8 |
| State 12: M=16 |
| Done! Computation length = 12. |

Figure 34: Execution of formula (7)

9.1.1 Incorporating temporal projection in Tempura

We add projection to Tempura by permitting statements of the form

$$w_1 \text{ proj } w_2.$$

Here w_1 and w_2 are themselves statements. Let us now look at two Tempura programs containing temporal projection.

Example (Specifying behavior in intermediate states):

The following formula initializes the variable M to 1 and then doubles it in every third state for 4 times:

$$([\text{len}(2) \wedge \text{stable } M]; \text{skip}) \text{ proj } [\text{len}(4) \wedge (M = 1) \wedge (M \text{ gets } 2M)]. \quad (7)$$

Furthermore M remains stable during intermediate states of the projection. Figure 34 gives an execution depicting the behavior of M .

| |
|--------------------------------|
| State 0: I= 0 J= 0 |
| State 1: I= 1 J= 0 |
| State 2: I= 1 J= 0 |
| State 3: I= 2 J= 1 |
| State 4: I= 2 J= 1 |
| State 5: I= 2 J= 1 |
| State 6: I= 3 J= 3 |
| State 7: I= 3 J= 3 |
| State 8: I= 3 J= 3 |
| State 9: I= 3 J= 3 |
| State 10: I= 4 J= 6 |
| Done! Computation length = 10. |

Figure 35: Execution of formula (8)

Example (Variable-length projection):

Below is a formula that initializes the variables I and J to 0 and then increases I by 1 and J by I with ever widening gaps dependent on I 's values:

$$\begin{aligned} & ([len(I) \wedge stable I \wedge stable J]; skip) \\ & \quad proj [len(4) \wedge (I = 0) \wedge (J = 0) \wedge (I \text{ gets } I + 1) \wedge (J \text{ gets } J + I)]. \end{aligned} \quad (8)$$

The values of I and J in intermediate states remain stable. This formula is executed in figure 35. Note that if we are not interested in the behavior of I and J in intermediate states, the projection can be specified as follows:

$$len(I + 1) \quad proj [len(4) \wedge (I = 0) \wedge (J = 0) \wedge (I \text{ gets } I + 1) \wedge (J \text{ gets } J + I)].$$

9.1.2 Universal projection

The kind of projection so far presented can be called *existential projection* since it is true on an interval if there is *at least* one way to break up the interval. The operator

$$w_1 \text{ uproj } w_2$$

is called *universal projection* and is true on an interval σ if for *every way* we break up σ using w_1 , the projected result is satisfied by w_2 . We can express *uproj* as the

dual of *proj*:

$$w_1 \text{ uproj } w_2 \equiv_{\text{def}} \neg(w_1 \text{ proj } \neg w_2).$$

At present, we do not see how to include universal projection in Tempura.

Example (Projection of a clocked system):

Universal projection provides a way of abstracting from digital behavior involving an explicit clock signal to behavior at the register-transfer level. As an example of this, let us consider a simple system driven by clock pulses. We first define the pulse operator $\uparrow\downarrow \text{Clock}$ to be true on intervals where the bit signal *Clock* rises and then falls:

$$\uparrow\downarrow \text{Clock} \equiv_{\text{def}} (\text{Clock} \approx 0); \text{skip}; (\text{Clock} \approx 1); \text{skip}; (\text{Clock} \approx 0).$$

The system we have in mind has the three bit signals *Clock*, *X* and *Y*. Universal projection provides a way of examining those properties of *X* and *Y* that are true across clock pulses independent of where we mark the beginning and end of each individual pulse. We can use projection of the form

$$(\uparrow\downarrow \text{Clock}) \text{ uproj } w$$

to specify and reasoning about a formula *w* describing the behavior of *X* and *Y* at the register-transfer level. For example, the following formula is true if *X* is repeatedly inverted and *Y* remains stable over a series of clock pulses:

$$(\uparrow\downarrow \text{Clock}) \text{ uproj } [(X \text{ gets } \neg X) \wedge \text{stable } Y].$$

9.2 Lambda Expressions and Pointers

Lambda expressions provide a natural means of representing functions as values. We now briefly sketch how to incorporate them into ITL and Tempura. We also show how lambda expressions can be used to represent pointers.

9.2.1 Lambda expressions

A lambda expression has the form

$$\lambda V_1, \dots, V_n: e,$$

where V_1, \dots, V_n are variables and e is an expression. The body of e can itself contain lambda expressions. In addition, we permit lambda predicates having the syntax

$$\Lambda V_1, \dots, V_n: w,$$

where V_1, \dots, V_n are variables and w is a formula.

First order lambda constructs can be viewed as temporal functions and predicates. To date, we do not have an adequate semantics for higher order lambda expressions.

An invocation of a lambda construct has the form

$$e_0(e_1, \dots, e_n),$$

where $n \geq 0$ and e_0, \dots, e_n are expressions. The value of e_0 should be a lambda construct of arity n .

9.2.2 Pointers

We can represent a pointer to a variable A by means of the parameterless lambda expression

$$\lambda: A.$$

For example, suppose the static variable b has this as its value. The variable A can then be indirectly accessed through the function invocation $b()$. Thus, the following formula is true on intervals where A increases by 1:

$$[b = (\lambda: A)] \wedge [b() \leftarrow b() + 1].$$

Let us now introduce some operators to make the pointer notation look more conventional. The expression

$$ref(l)$$

equals a pointer to the location l . It can be defined as follows:

$$ref(l) \quad =_{\text{def}} \quad \lambda: l.$$

A pointer expression e can be dereferenced using the construct

$$deref(e).$$

It is defined as follows:

$$\mathit{deref}(e) \equiv_{\text{def}} e().$$

Thus the sample formula given above can be also expressed as follows:

$$[b = \mathit{ref}(A)] \wedge [\mathit{deref}(b) \leftarrow \mathit{deref}(b) + 1].$$

9.3 The *process* Construct

In section 8.9.6, we mentioned that two separate statements can not in parallel determine the length of an interval. If one has a number of statements w_1, w_2, \dots, w_n to be run in parallel and each individually determines interval length, then all but one should be used with the unary operator *process*:

$$w_1 \wedge (\mathit{process} w_2) \wedge \dots \wedge (\mathit{process} w_n).$$

For instance, the following statement is permitted:

$$\begin{aligned} & [(I = 0) \wedge (I \text{ gets } I + 1) \wedge \mathit{halt}(I = 5)] \\ & \wedge \mathit{process} [(J = 0) \wedge (J \text{ gets } J + 2) \wedge \mathit{halt}(J = 10)]. \end{aligned}$$

The predicate *sum_tree_process* in section 7.1.2 and the predicate *quicksort_process* in section 7.3.3 both make use of *process*. This is so that multiple instances of them can be run in parallel. Note that the *process* construct has no logical semantics and can be defined as follows:

$$\mathit{process} w \equiv_{\text{def}} w.$$

Its purpose is merely to explicitly indicate to the interpreter that redundant specifications of interval termination are present.

9.4 Implementing these Constructs

We now discuss how projection, lambda expressions, pointers and the *process* construct can be implemented in the Tempura interpreter. This presentation can be skipped if desired.

9.4.1 Implementing temporal projection

The interpreter handles a statement of the form $w_1 \text{ proj } w_2$ by first allocating a new memory cell and then transforming the statement to the internal construct

$$\text{@proj}(\text{done-cell}, w_1, w_2).$$

This is immediately re-reduced. Here *done-cell* is the index of the memory cell. The cell serves as a local done-flag for the projected interval in which the statement w_2 is executed.

We execute the *@proj* construct by first saving the value of the variable *Current_Done_Cell* and setting it to the index *done-cell*. The statement w_2 is then transformed in this context to a new statement w' . Afterwards the old value of *Current_Done_Cell* is restored. If w' is not yet fully reduced, the overall *@proj* statement is rewritten as

$$\text{@proj}(\text{done-cell}, w_1, w').$$

This is returned as the result of the transformation. On the other hand, if w' is fully reduced and thus of the form $\odot w''$, then the overall projection is transformed to the following conditional statement and then immediately re-reduced:

$$\text{if } @loc(\text{done-cell}) \text{ then empty else } [w_1; \text{@proj}(\text{done-cell}, w_1, w'')].$$

This tests the localized done-flag indexed by *done-cell*. If it is true, the interval in which w_2 was transformed is finished and therefore the overall projection terminates. Otherwise, the statement w_1 is executed followed by the resumption of the projection statement.

9.4.2 Implementing lambda expressions and pointers

The lambda constructs implemented in the Tempura interpreter can uniformly handle both first-order and higher-order variants. The approach taken is basically the same as for processing predicate and function definitions. See section 8.7.10 for a presentation of this.

The technique of using lambda expressions for representing pointers actually works in the current Tempura interpreter. However, the constructs *ref* and *deref* have been implemented in a more efficient but less general manner by means of a new type of descriptor specifically for pointers. We omit the details.

9.4.3 Implementing the *process* construct

When reducing a statement of the form

process w

we first allocate a memory cell to serve as a local done-flag. The statement is then transformed to the internal construct

@process(done-cell, w)

and immediately re-reduced. Here *done-cell* is the index of the memory cell.

The *@process* construct is itself implemented by saving the value of the variable *Current_Done_Cell* and setting it to the index *done-cell*. The statement *w* is transformed within this context to a new statement *w'*. Afterwards the variable *Current_Done_Cell* is restored to its old value.

If the statement *w'* is not yet fully reduced, the statement

@process(done-cell, w')

is returned as the result. Otherwise, if *w'* is fully reduced then it has the form $\odot w''$ for some *w''*. We therefore transform the *@process* construct to the following conditional statement:

if [empty \equiv @loc(done-cell)] then [\odot @process(done-cell, w'')]
else false.

This tests to make sure that the current done-flag and the local done-flag used by the *@process* operator agree in value. If they do, the statement *w''* is placed within the *@process* construct in preparation for any subsequent states. If the done-flags are not equal, the statement *false* is executed, thus generating an error. Note that once a Tempura program is debugged, the consistency check between the done-flags can be suppressed.

10 Experience and Further Work

Using the ideas discussed here, we have implemented a prototype Tempura interpreter in Lisp. Its design is based on the interpreter presented in section 8

and it incorporates some of the modifications mentioned in section 8.9. A great variety of Tempura programs have been written and successfully run. Roger Hale, a PhD student at Cambridge University, has more recently implemented a faster version of the interpreter in the programming language *C*. In [7], he describes the application of ITL and Tempura to the modelling of a digital ring network. Another interpreter has been developed in Prolog by Masahiro Fujita, Shinji Kono and others at Tokyo University.

In the future we plan to build a compiler and bootstrap Tempura in itself. We also hope to describe the operational semantics of Tempura in ITL. This will enable us to formalize the relation between various execution strategies, both sequential and parallel.

So far we have made little mention of ITL's proof theory. This is much less developed than the model theory. Work by Halpern and Moszkowski has shown that one propositional subset of ITL is undecidable and that another is decidable (see Moszkowski [25] for details). Although it is not hard to come up with sound axioms and inference rules for ITL, no systematic work has been done to date. We therefore feel that ITL's proof theory represents a promising area for future research.

11 Some Related Work

We now look at some logical formalisms and programming languages that seem related to ITL and Tempura.

11.1 Other Temporal Logics Based On Intervals

A number of researchers have proposed variants of temporal logic that facilitate reasoning about intervals. Schwartz, Melliar-Smith and Vogt [31] develop a formalism called *interval logic* which includes temporal formulas having the syntax

$$[I]w.$$

Here *I* can be built from a variety of special constructs for indicating the scope of the interval in which the temporal formula *w* is to be evaluated. Thus *I* can be thought of as an "interval designator." For example, the formula

$$[(X = Y) \Rightarrow (Y = 16)] \square (X > Z)$$

is true on an interval if X is greater than Z throughout the subinterval starting the first time X equals Y and ending when Y equals 16. Note that the designator is not itself a temporal formula. This is unlike the approach of ITL in which constructs used to specify subintervals are themselves always formulas.

Shasha, Pnueli and Ewald [33] propose some generalized forms of the *next* operator in which explicit time offsets are mentioned. For example, the formula

$$\bigcirc^e w$$

is true if the subformula w is true in e units of time from now. This is basically the same as the ITL formula

$$\text{len}(e); w.$$

The following two types of formulas are also permitted:

$$[e_1, e_2]w, \quad \langle e_1, e_2 \rangle w.$$

They are defined as follows:

$$\begin{aligned} [e_1, e_2]w &\equiv_{\text{def}} \forall i: ((e_1 \leq i < e_2) \supset \bigcirc^i w), \\ \langle e_1, e_2 \rangle w &\equiv_{\text{def}} \exists i: ((e_1 \leq i < e_2) \wedge \bigcirc^i w). \end{aligned}$$

These constructs let one express behavior over various suffix subintervals of time but do not provide access to prefix subintervals in the way ITL's *chop* does. For example, the following ITL formula does not seem to be as elegantly expressed in their notation:

$$(K + 1 \rightarrow K); (K + 2 \rightarrow K)$$

11.2 The Programming Language *Lucid*

The functional programming language *Lucid* [2,35] developed by Ashcroft and Wadge is similar to parts of Tempura. For example, the *Lucid* program

$$I = 0 \text{ fby } (I + 1); \quad J = 0 \text{ fby } (J + I)$$

roughly corresponds to the temporal formula

$$(I = 0 \wedge J = 0) \wedge (I \text{ gets } I + 1) \wedge (J \text{ gets } J + I).$$

This illustrates how the operator *gets* can be handled in Lucid. On the other hand, Algol-like temporal constructs such as \leftarrow , *chop* and *while* do not have direct analogs in Lucid. Thus, a Tempura construct such as

$$\textit{while } (M \neq 0) \textit{ do } (\textit{skip} \wedge [M \leftarrow M - 1] \wedge [N \leftarrow 2N])$$

cannot be readily translated. In [1], Ashcroft and Wadge develop a calculus for reasoning about Lucid programs.

11.3 CCS and CSP

Milner's *Calculus of Communicating Systems* [23] and Hoare's *Communicating Sequential Processes* [15] are both popular notations for describing and reasoning about parallel systems. Related work includes Milne's *CIRCAL* [22] for modelling hardware and the CSP-inspired programming language *occamTM* [17]. In CCS and CSP, multiple processes interact with one another by mutually synchronizing on events. There are operators for composing processes, waiting for events and concealing events. CSP includes constructs for representing channels.

The treatment of time is not a central issue in these formalisms. However, it is possible to describe a clock process that serves as a source of events representing ticks. A variant of CCS called *Synchronous CCS* models concurrent systems that operate in lock step.

11.4 Predicative Programming

Hehner [11] views programs as logical predicates over input-output behavior of variables. Various Algol-based constructs such as assignment (" $:=$ "), sequencing (" $;$ ") and while-loops are treated in this manner. Time is introduced by means of an explicit clock variable. Hehner then goes on to introduce concurrency through CSP-like channels. It seems too early to compare this approach with temporal logic.

11.5 Synthesis and Verification of Specifications

Manna and Wolper [21] investigate techniques for synthesizing CSP synchronization code from temporal logic specifications. Examples include a reader-writer

system and the dining philosophers' problem. Mishra and Clarke [24] use temporal logic to describe asynchronous digital circuits and then generate corresponding finite-state automata. Tang [34] and Manna and Pnueli [20] discuss ways of translating conventional programming constructs into transition systems described in temporal logic. The resulting temporal descriptions can then be used to reasoning about the original programs. In contrast to these approaches, we have worked on developing a set of temporal operators that facilitate programming directly in temporal logic. This bypasses the need for two notations and omits any synthesis from specifications.

11.6 The Programming Language *Prolog*

The programming language *Prolog* [5,16] uses Horn clauses in first order logic as a means of expressing computations. The original applications for this approach were in areas of artificial intelligence such as natural language understanding and theorem proving. This led to the use of resolution as the basis for executing Prolog programs. Subsequent work by Shapiro [32], Clark and Gregory [4] and others has generalized logic programming to deal with concurrent systems.

Prolog and its offspring do not have any notion of time and therefore are unable to directly express imperative constructs such as assignments. In practice, extra-logical constructs such as *assert* and *retract* can be used to get around this restriction. They provide a means for adding and removing facts from a system-maintained database. However, their usage is generally not considered good programming style. Alternatively, one can stay completely within the framework of the underlying logic by explicitly representing dynamically changing objects as lists of values or by including times as parameters to predicates.

In comparison with Prolog, our approach does not limit itself to a subset of convention logic based on Horn clauses and does not use resolution as the means of executing formulas. Nonetheless, it is a form of logic programming, albeit with temporal logic as the underlying formalism. We feel this provides a more natural setting for both programming and reasoning about dynamic systems.

11.7 Functional Programming

Functional programming [12] is based on the idea that certain types of functions can be interpreted as executable descriptions of computations. McCarthy's

programming language *Lisp* [18] is perhaps the best known example of this approach. The functions themselves have no built-in notion of time. Therefore, as is the case with logic programs, dynamic behavior is in effect modelled indirectly. Recursion seems to be the most common technique used for this. Another approach is to represent time-dependent variables as lists or as functions containing an explicit time parameter.

Constructs such as Lisp's *setq* and *rplaca* are sometimes included in an ad hoc manner in order to permit in-place assignments and other imperative operations. Language purists tend to discourage their use. Even so, such features often seem indispensable for reasons of clarity and efficiency.

12 Conclusions

The present work has investigated Tempura, a program language based on Interval Temporal Logic. The ITL formalism provides a way to treat such programming concepts as assignment and loops as formulas about intervals of time. Therefore, Tempura programs, their specifications and their properties can all be expressed in the same formalism. Furthermore, this approach provides a unified way for modelling both hardware and software. In the future, we hope to gain more experience with using ITL and Tempura to simulate and reason about descriptions of hardware devices and other types of parallel systems. In addition, we plan to explore the feasibility of using Tempura as a general-purpose programming language.

Acknowledgements

Roger Hale constructed a parser for Tempura and patiently served as the resident guinea pig and captive user during the development of the system. In addition, he critically read various drafts of this report. We wish to thank him as well as Bill Clocksin, Mike Gordon and Edmund Ronald for their stimulating conversations and suggestions. Funding from the British Science and Engineering Research Council is gratefully acknowledged.

References

- [1] E. A. Ashcroft and W. W. Wadge. Lucid: A formal system for writing and proving programs. *SIAM Journal of Computing* 5, 3 (September 1976), 336-354.
- [2] E. A. Ashcroft and W. W. Wadge. Lucid, a nonprocedural language with iteration. *Communications of the ACM* 20, 7 (July 1977), 519-526.
- [3] A. Chandra, J. Halpern, A. Meyer, and R. Parikh. Equations between regular terms and an application to process logic. *Proceedings of the 13-th Annual ACM Symposium on Theory of Computing*, Milwaukee, Wisconsin, May, 1981, pages 384-390.
- [4] K. L. Clark and S. Gregory. Parlog: Parallel programming in logic. Technical Report DOC 84/4, Department of Computing, Imperial College, London, April, 1984. To appear in *ACM Transactions on Programming Languages and Systems*.
- [5] W. F. Clocksin and C. S. Mellish. *Programming in Prolog*. Springer-Verlag, Berlin, 1981.
- [6] M. J. C. Gordon. LCF-LSM: A system for specifying and verifying hardware. Computer Laboratory, University of Cambridge, Technical Report 41, 1983.
- [7] R. W. S. Hale. Modelling a ring network in interval temporal logic. To appear in *Proceedings of Euromicro 85*, Brussels, Belgium, September, 1985.
- [8] J. Halpern, Z. Manna and B. Moszkowski. A hardware semantics based on temporal intervals. In *Proceedings of the 10-th International Colloquium on Automata, Languages and Programming*, Barcelona, Spain, July, 1983.
- [9] D. Harel. *First-Order Dynamic Logic*. Number 68 in the series *Lecture Notes in Computer Science*, Springer-Verlag, Berlin, 1979.
- [10] D. Harel, D. Kozen, and R. Parikh. Process logic: Expressiveness, decidability, completeness. *Journal of Computer and System Sciences* 25, 2 (October 1982), pages 144-170.
- [11] E. C. R. Hehner. Predicative programming (parts I and II). *Communications of the ACM* 27, 2 (February 1984), pages 134-151.

- [12] P. Henderson. *Functional Programming: Application and Implementation*. Prentice-Hall International, London, 1980.
- [13] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM* 12, 10 (October 1969), pages 576–580, 583.
- [14] C. A. R. Hoare. Towards a theory of parallel programming. In C. A. R. Hoare and R. H. Perrott, editors, *Operating Systems Techniques*, pages 61–71. Academic Press, London, 1972.
- [15] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice Hall International, London, 1985.
- [16] C. J. Hogger. *Introduction to Logic Programming*. Academic Press, Orlando, Florida, 1984.
- [17] Inmos Ltd. *OccamTM Programming Manual*. Prentice Hall International, London, 1984.
- [18] J. McCarthy, P. W. Abrahams, D. J. Edwards et al. *Lisp 1.5 Programmer's Manual*. MIT Press, Cambridge, Massachusetts, 1962.
- [19] Z. Manna and A. Pnueli. Verification of concurrent programs: The temporal framework. In R. S. Boyer and J. S. Moore, editors, *The Correctness Problem in Computer Science*, pages 215–273, Academic Press, New York, 1981.
- [20] Z. Manna and A. Pnueli. How to cook your favorite programming language in temporal logic. In *Proceedings of the Tenth Annual ACM Symposium on Principles of Programming Languages*, Austin, Texas, January, 1983, pages 141–154.
- [21] Z. Manna and P. L. Wolper. Synthesis of computing processes from temporal logic specifications. *ACM Transactions on Programming Languages and Systems* 6, 1 (January 1984), pages 68–93.
- [22] G. J. Milne. CIRCAL and the representation of communication, concurrency and time. *ACM Transactions on Programming Languages and Systems* 7, 2 (April 1985), pages 270–298.
- [23] R. Milner. *A Calculus of Communicating Systems*. Number 92 in the series *Lecture Notes in Computer Science*, Springer-Verlag, Berlin, 1980.

- [24] B. Mishra and E. M. Clarke, Automatic and hierarchical verification of asynchronous circuits using temporal logic. Technical report CMU-CS-83-155, Department of Computer Science, Carnegie-Mellon University, September, 1983.
- [25] B. Moszkowski. *Reasoning about Digital Circuits*. PhD Thesis, Department of Computer Science, Stanford University, 1983.
- [26] B. Moszkowski. A temporal analysis of some concurrent systems. To appear in *Proceedings of the STL Workshop on Concurrency*, Cambridge, England, September, 1983.
- [27] B. Moszkowski. A temporal logic for multilevel reasoning about hardware. *Computer* 18, 2 (February 1985), pages 10–19.
- [28] B. Moszkowski and Z. Manna. Reasoning in interval temporal logic. Technical report STAN-CS-83-969, Department of Computer Science, Stanford University, July, 1983.
- [29] V. R. Pratt, Semantical considerations on Floyd-Hoare logic. In *Proceedings of the 17-th Annual IEEE Symposium on Foundations of Computer Science*, Houston, Texas, October, 1976, pages 109–121.
- [30] N. Rescher and A. Urquhart. *Temporal Logic*. Springer-Verlag, New York, 1971.
- [31] R. L. Schwartz, P. M. Melliar-Smith, and F. H. Vogt. An interval logic for higher-level temporal reasoning. In *Proceedings of the Second Annual ACM Symposium on Principles of Distributed Computing*, Montreal, Quebec, Canada, August, 1983, pages 173–186.
- [32] E. Y. Shapiro. Systems programming in Concurrent Prolog, In *Proceedings of the Eleventh Annual ACM Symposium on Principles of Programming Languages*, Salt Lake City, Utah, January, 1984, pages 93–105.
- [33] D. E. Shasha, A. Pnueli and W. Ewald. Temporal verification of carrier-sense local area network protocols. In *Proceedings of the Eleventh Annual ACM Symposium on Principles of Programming Languages*, Salt Lake City, Utah, January, 1984, pages 54–65.

- [34] C. Tang. Toward a unified logic basis for programming languages. In *Proceedings of IFIP Congress 83*, Elsevier Science Publishers B.V. (North-Holland), Amsterdam, 1983, pages 425–429.
- [35] W. W. Wadge and E. A. Ashcroft. *Lucid, the Dataflow Programming Language*. Academic Press, Orlando, Florida, 1985.