

Specifying Concurrent Objects as Communicating Processes*

Jayadev Misra

Department of Computer Sciences

The University of Texas at Austin

Austin, Texas 78712

(512) 471-9547

`misra@cs.utexas.edu`

1 Introduction

An object and its associated operations may be specified in many ways. One way is to give an abstract representation of the object data structure (viz., representing a queue by a sequence) and the effects of various operations on this abstract representation [3,4,5]. Another way [2] is to leave the representation aspects unspecified but to give a set of equations that relate the effects of various operations (the equations define a congruence relation in a term algebra). The common goal of a specification, however, is to serve as a legal document that spells out a contract between a user and an implementer: The user may assume no more about an object than what is specified and the implementer must satisfy the specification through his implementation.

We view a specification not merely as a legal contract but additionally as a means (1) to deduce properties of a specified object, (2) to deduce properties of other objects in which the specified object is a component (i.e., inheritance of properties) and (3) to implement the object by stepwise refinement of its specification. Therefore we require that a specification not merely be formal but also be in a form that admits of effective manipulation. This requirement rules out specification schemes in which program fragments (in some high level language) appear as part of specifications; typically such program fragments cannot be manipulated effectively.

In many specification schemes it is assumed that (1) each operation on an object is deterministic (i.e., applying the operation to a given state of the object results in a unique next state and/or unique values being returned), (2) an operation once started always terminates in every state of the object, and (3) operations are not applied concurrently. In many cases of interest arising in applications such as operating systems, process

*This material is based in part on work supported by the Texas Advanced Research Program under Grant No. 003658-065 and by the ONR Contract N00014-90-J-1640.

control systems and concurrent databases, these assumptions are rarely met. For instance, a “queue object” acts as an intermediary between a producer and a consumer, temporarily storing the data items output by the producer and later delivering them to the consumer. The queue object is required to (1) deliver data in the same order to the consumer as they were received from the producer, (2) receive data from the producer provided its internal queue spaces are nonfull, and (3) upon demand, send data to the consumer provided its internal queue spaces are nonempty. Requests from the producer and the consumer may be processed concurrently: The producer is delayed until there is some space in the queue and the consumer is delayed until there is some data item in the queue. Observe that a request from the producer, to add a data item to the queue, may not terminate if the queue remains full forever, and similarly, a request from the consumer may not terminate if the queue remains empty forever.

In this paper, we propose a specification scheme for concurrent objects that allows effective manipulations of specifications and admits nondeterministic, nonterminating and concurrent operations on objects. Our approach is to view a concurrent object as an asynchronous communicating process. Such a process can be specified by describing its initial state, how each communication—send or receive—alters the state, and the conditions under which a communication action is guaranteed to take place. On the other hand, since the internal state can be determined from the sequence of communications (provided the process is deterministic), the process can also be specified in terms of the sequence of communications in which it engages. The point to note is that the internal state and the sequence of communications can be viewed as *auxiliary variables* which may be altered as a result of communications. Our specification scheme allows us to define auxiliary variables and state properties using these variables. We advocate using any auxiliary variable that allows a simple and manipulable specification, leaving open the question whether it is preferable to use “observable” input-output sequences or “unobservable” internal states.

The specification mechanism is based on “UNITY logic” [1]. Auxiliary variables can be defined directly using the constructs in this logic. The inference rules of the logic can be applied to deduce various properties of an object from its specification. Also, we may use the logic to prove that one specification refines (i.e., implements) another specification and that an implementation meets a specification.

1.1 Outline of the Paper

We give a brief introduction to UNITY in Section 2 including all the notations and logic to understand this paper. The main example treated in this paper is a bounded bag (or multiset). Section 3 contains an informal description and a formal specification of this object. We demonstrate the usefulness of the specification by showing that concatenation of two bags results in a bag of the appropriate size. The specification in Section 3 is based on an auxiliary variable that encodes the internal state. We provide an alternative specification in Section 4 that uses the sequences of input/output as auxiliary variables; we show that this specification is a

refinement of the earlier one. Section 5 contains specifications for a “queue” and a “stack.” We show that a queue implements a bag and a stack implements a bag, showing in each case that the specification of the former implies the specification of the latter. Section 6 contains a refinement of the bag specifications of Section 4, and Section 7 gives an implementation of the specification of Section 6. The usefulness of the method is discussed in Section 8.

All the proofs are given in complete detail to emphasize that such proofs need not be excessively long or tedious, as is often the case with formal proofs.

A preliminary version of this paper appears in [7]. The specifications and the proofs have been considerably simplified in the current version.

2 A Brief Introduction to UNITY

A UNITY program consists of (1) declarations of variables, (2) a description of their initial values and (3) a finite set of statements. We shall not describe the program syntax except briefly in Section 7 of this paper because it is unnecessary for understanding this paper. However an operational description of the execution of a program is helpful in understanding the logical operators introduced later in this section.

An initial state of a program is a state in which all variables have their specified initial values (there can be several possible initial states if the initial values of some variables are not specified). In a step of program execution an arbitrary statement is selected and executed. Execution of every statement terminates in every program state. (This assumption is met in our model by restricting the statements to assignment statements only, and function calls, if any, must be guaranteed to terminate.) A program execution consists of an infinite number of steps in which each statement is executed infinitely often.

This program model captures many notions useful for programming such as: synchrony, by allowing several variable values to be modified by a single (atomic) statement; asynchrony, by specifying little about the order in which the individual statements are executed; processes communicating through shared variables, by partitioning the statements of the program into subsets and identifying a process with a subset; processes communicating via messages, by restricting the manner in which shared variables are accessed and modified, etc. We shall not describe these aspects of the model; however, it will become apparent from the bag example in this paper that process networks can be described effectively within this model.

2.1 UNITY logic

A fragment of UNITY logic, that is used in this paper, is described in this section.

Three logical operators, *unless*, *ensures*, and *leads-to*, are at the core of UNITY logic. Each of these is a binary relation on predicates; *unless* is used to describe the safety properties, and the other two to describe

progress properties of a program.

Notation: Throughout this section p, q, r denote predicates which may name program variables, bound variables and free variables (free variables are those that are neither program variables nor bound variables). \square

2.1.1 unless

For a given program

$$p \text{ unless } q$$

denotes that once predicate p is true it remains true at least as long as q is not true. Formally (using Hoare's notation)

$$\frac{\langle \text{for all statements } s \text{ of the program} \ :: \ \{p \wedge \neg q\} \ s \ \{p \vee q\} \rangle}{p \text{ unless } q}$$

i.e., if $p \wedge \neg q$ holds prior to execution of any statement of the program then $p \vee q$ holds following the execution of the statement.

It follows from this definition that if p holds and q does not hold in a program state then in the next state either q holds, or $p \wedge \neg q$ holds; hence, by induction on the number of statement executions, $p \wedge \neg q$ continues to hold as long as q does not hold. Note that it is possible for $p \wedge \neg q$ to hold forever. Also note that if $\neg p \vee q$ holds in a state then $p \text{ unless } q$ does not tell us anything about future states.

Notation: We write $\langle \forall u \ :: \ P.u \rangle$ and $\langle \exists u \ :: \ P.u \rangle$ for universal quantification of u in $P.u$ and existential quantification of u in $P.u$, respectively. The dummy u could denote a variable, a statement in a program, or even a program. Any property having a free variable (i.e., a variable that is neither bound nor a program variable) is assumed to be universally quantified over all possible values of the variable. Thus,

$$u = k \text{ unless } u > k$$

where k is free, is a shorthand for

$$\langle \forall k \ :: \ u = k \text{ unless } u > k \rangle .$$

\square

Examples

1. Integer variable u does not decrease.

$$u = k \text{ unless } u > k$$

or, $u \geq k$ unless false

2. A message is received (i.e., predicate *rcvd* holds) only if it had been sent earlier (i.e., predicate *sent* already holds).

$\neg rcvd$ unless *sent*

□

Example (Defining Auxiliary Variables)

Let u be an integer-valued variable and let v count the number of times u 's value has been changed during the course of a program execution. The value of v is completely defined by

initially $v = 0$

$u = m \wedge v = n$ unless $u \neq m \wedge v = n + 1$

The traditional way to define v , given a program in which u is a variable, is to augment the program text with the assignment,

$v := v + 1$

whenever u 's value is changed; v is called an auxiliary variable. Our way of defining v , without appealing to the program text, is preferable because it provides a direct relationship between u and v which may be exploited in specifications and proofs. □

2.1.2 Stable, Constant, Invariant

Some special cases of *unless* are of importance. The predicate p unless false, from definition, denotes that p remains true forever once it becomes true; we write, “ p stable” as a shorthand for “ p unless false.” An expression e is *constant* means $e = x$ is stable, for all possible values x of e ; then e never changes value. If p holds in every initial state and p is stable then p holds in every state during any execution; we then say that “ p is invariant.”

2.1.3 ensures

For a given program,

p ensures q

implies that p unless q holds for the program and if p holds at any point in the execution of the program then q holds eventually. Formally,

$$\frac{p \text{ unless } q \wedge \langle \exists \text{ statement } s :: \{p \wedge \neg q\} s \{q\} \rangle}{p \text{ ensures } q}$$

It follows from this definition that once p is true it remains true at least as long as q is not true (from p unless q). Furthermore, from the rules of program execution, statement s will be executed sometime after p becomes true. If q is still false prior to the execution of s then $p \wedge \neg q$ holds and the execution of s establishes q .

2.1.4 leads-to

For a given program, p leads-to q , abbreviated as $p \mapsto q$, denotes that once p is true, q is or becomes true. Unlike *ensures*, p may not remain true until q becomes true. The relation *leads-to* is defined inductively by the following rules. The first rule is the basis for the inductive definition of *leads-to*. The second rule states that *leads-to* is a transitive relation. In the third rule, $p.m$, for different m , denote a set of predicates. This rule states that if every predicate in a set *leads-to* q then their disjunction also *leads-to* q .

(basis)

$$\frac{p \text{ ensures } q}{p \mapsto q}$$

(transitivity)

$$\frac{p \mapsto q, q \mapsto r}{p \mapsto r}$$

(disjunction) In the following m ranges over any arbitrary set, and m is not free in q .

$$\frac{\langle \forall m :: p.m \mapsto q \rangle}{\langle \exists m :: p.m \rangle \mapsto q}$$

Notes on Inference Rules

We have explained the meaning of each logical operator in terms of program execution. However, neither the definitions nor our proofs make any mention of program execution. We use only the definitions, and a few rules derived from these definitions, in proofs; we believe that our proofs are succinct because we avoid operational arguments about program executions.

2.1.5 Derived Rules

The following rules for *unless* are used in this paper; for their proofs, see [1].

(reflexivity, antireflexivity)

$$p \text{ unless } p, p \text{ unless } \neg p$$

(consequence weakening)

$$\frac{p \text{ unless } q, q \Rightarrow r}{p \text{ unless } r}$$

(stable conjunction)

$$\frac{p \text{ unless } q, r \text{ stable}}{p \wedge r \text{ unless } q \wedge r}$$

We need the following facts about constants; see [6] for proofs.

(constant formation)

Any expression built out of constants and free variables is a constant.

(constant introduction) For any function f over program variables u ,

$$\frac{u = k \text{ unless } u \neq k \wedge f(u) = f(k)}{f \text{ constant}}$$

Corollary 1: For any predicate p over program variables u ,

$$\frac{u = k \text{ unless } u \neq k \wedge p}{p \text{ stable}}$$

We use the following results about *leads-to*.

(Implication)

$$\frac{p \Rightarrow q}{p \mapsto q}$$

The following rule allows us to deduce a progress property from another progress property and a safety property.

(PSP)

$$\frac{p \mapsto q, r \text{ unless } b}{p \wedge r \mapsto (q \wedge r) \vee b}$$

2.1.6 Substitution Axiom

Substitution axiom allows us to replace an invariant by “true” and vice versa, in any predicate. Thus, if I is an invariant and it is required to prove that

$$p \mapsto q \wedge I$$

it suffices to prove

$$p \mapsto q .$$

It is important to note that when several programs are composed (see Section 2.2), the substitution axiom can be applied only with an invariant of the composite program.

2.2 Program Composition through *union*

Given two programs F, G , their *union*, written $F \parallel G$, is obtained by appending their codes together: The initial conditions of both F, G are satisfied in $F \parallel G$ (and hence, we assume that initial conditions of F, G are not contradictory) and the set of statements of $F \parallel G$ is the union of the statements of F and G .

Programs F, G may be thought of as executing asynchronously in $F \parallel G$. The union operator is useful for understanding process networks where each process may be viewed as a program and the entire network is their union; in this view, the bag program, the producer, and the consumer are composed through the union operator to form a system.

The following theorem is fundamental for understanding union. It says that an *unless* property holds in $F \parallel G$ iff it holds in both F and G ; an *ensures* property holds in $F \parallel G$ iff the corresponding *unless* property holds in both components and the *ensures* property holds in at least one component. (When there are multiple programs we write the program name with a property, such as $p \text{ unless } q \text{ in } F$.)

Union Theorem:

$$\begin{aligned} p \text{ unless } q \text{ in } F \parallel G &\equiv p \text{ unless } q \text{ in } F \wedge p \text{ unless } q \text{ in } G \\ p \text{ ensures } q \text{ in } F \parallel G &\equiv (p \text{ unless } q \text{ in } F \wedge p \text{ ensures } q \text{ in } G) \\ &\quad \vee (p \text{ ensures } q \text{ in } F \wedge p \text{ unless } q \text{ in } G) \end{aligned}$$

Corollary 1:

$$\frac{p \text{ unless } q \text{ in } F, p \text{ stable in } G}{p \text{ unless } q \text{ in } F \parallel G}$$

Corollary 2:

$$\frac{e \text{ constant in } F, e \text{ constant in } G}{e \text{ constant in } F \parallel G}$$

Corollary 3:

$$\frac{p \text{ stable in } F, p \text{ invariant in } G}{p \text{ invariant in } F \parallel G}$$

Corollary 4:

$$\frac{p \text{ ensures } q \text{ in } F, p \text{ stable in } G}{p \text{ ensures } q \text{ in } F \parallel G}$$

If variable x cannot be accessed by a program F —i.e., x is local to some other program—then x is constant in F ; hence—using the constant formation rule—any expression, $e(x)$, over x is constant in F . This fact is expressed in Corollary 5. We write “ x is not accessed in F ” in quotes because this is not a proposition in our logic.

Corollary 5:

$$\frac{\text{“}x \text{ is not accessed in } F\text{”}}{e(x) \text{ constant in } F}$$

We shall also use properties of the following form in the specification of F , where P, Q are arbitrary properties. In the following, G is quantified over programs:

$$\langle \forall G :: (P \text{ in } G) \Rightarrow (Q \text{ in } F \parallel G) \rangle$$

This says that if P is a property of any program G then Q is a property of $F \parallel G$. This is a convenient way of specifying the properties of the environment with which F may be composed and the resulting properties of the composite program. This kind of formula will be crucial in specifications of concurrent objects, because the objects typically assume certain properties of their environments.

3 Specification Methodology; Bag Specification

3.1 Specification Methodology

We treat a concurrent object as an asynchronous communicating process. A process and its environment communicate over channels that can typically hold several items of data. To simplify matters, we dispense with channels; we assume that communications are through certain shared variables each of which can hold at most one item of data. The access protocol to the shared variables is as follows: An “input shared variable” of the object can be written only by the environment and read by the object, an “output shared variable” of the object can be written only by the object and read by the environment. Typically, we will have a shared variable corresponding to each operation on the object (if an operation also delivers a result, we may require two shared variables—one to request the operation and the other to store the result).

In Figure 1, B implements a bag that is shared by a group of producers and consumers. Producers add items to the bag by successively storing them in r ; consumers remove successive items from w . Program F is the environment of B , representing the producers and the consumers. There might be multiple producers and




Figure 1: A Program B and its environment F

consumers or even a single process that is both the producer and the consumer; the exact number is irrelevant for the specification. The values that can be written into r, w are, again, irrelevant for specification. However we do postulate a special value, ϕ , which is written into a variable to denote that it is “empty,” i.e., it contains no useful data. The protocol for reading and writing is as follows. Program F writes into r only if $r = \phi$; program B reads a value from r only if $r \neq \phi$ and then it may set r to ϕ . Program B stores a value in w only if w is ϕ ; program F reads from w and it may set w to ϕ to indicate that it is ready to consume the next piece of data.

To formalize the access protocol of the last paragraph, we introduce an ordering relation, \prec , over the data values and ϕ as follows: ϕ is “smaller than” all non- ϕ values, i.e.,

$$X \prec Y \equiv X = \phi \wedge Y \neq \phi$$

Then it follows that

$$X \preceq Y \equiv X = \phi \vee X = Y$$

Property P1, below, states that B removes only non- ϕ data from r and it may set r to ϕ . Property P2 states that B writes only non- ϕ values into w provided w is ϕ . In the following, R, W are free variables (of type, data or ϕ).

P1. $r \preceq R$ stable in B

P2. $W \preceq w$ stable in B

The access protocol specification as given by P1 and P2 is identical for all concurrent objects. The initial conditions can usually be specified for output shared variables; for this example

P0. initially $w = \phi$ in B

Next, we specify the safety properties—the effect of reading inputs or how the outputs are related to the current state—and progress properties—the conditions under which inputs are read and outputs are written. These properties are specific to each object. They are explored next for bags.

3.2 Internal State as an Auxiliary Variable (An Informal Specification)

In the following specification we introduce an auxiliary variable, b , that represents the bag of data items internally stored by B . Variable b is local to B . The size of b does not exceed N , a given positive integer (Property P3). Property P4 states that any item read from r is added to b or stored in w (and r is then set to ϕ) and any item removed from b is stored in w (w was ϕ previous to this step); hence the union of r, b, w remains constant. Property P5 states that independent of the environment if the bag is nonfull ($|b| < N$) an item, if there is any, will be removed from r and analogously, if the bag is nonempty ($|b| > 0$) w is or will become non- ϕ .

Notation: All bags in this paper are finite. We treat r, w to be bags of size at most 1. Union of bags u, v is written as $u + v$; the bag $u - v$ is the bag obtained by removing as many items of v as possible from u . The value ϕ is treated as an empty bag.

Combining with P0, P1, P2 defined previously, we obtain

3.3 Specification of a Concurrent Bag of Size $N, N > 0$

There exists a variable b , local to B , of type bag such that

- P0. initially $b + w = \phi$ in B
- P1. $r \preceq R$ stable in B
- P2. $W \preceq w$ stable in B
- P3. $|b| \leq N$ invariant in B
- P4. $r + b + w$ constant in B
- P5. $\langle \forall F ::$
 - (P5.1) $|b| < N \mapsto r = \phi$ in $B \parallel F$
 - (P5.2) $|b| > 0 \mapsto w \neq \phi$ in $B \parallel F$

Note: A program that implements B may not have a variable named b . It is merely required that b be computable from the internal variables of B ; in this sense, b is an auxiliary variable of B .

Note: Since b is a local variable of B , it is not accessed by the environment F . Using Corollary 5 of Section 2.2, $|b| \leq N$ stable in F . Since $|b| \leq N$ is invariant in B , using Corollary 3 of Section 2.2,

- P6. $\langle \forall F :: |b| \leq N$ invariant in $B \parallel F \rangle$

Observe that P5 is a property of program B ; it says that if B is composed with any program F then certain properties hold in the composite program, $B \parallel F$. In particular, P5 does not require F to obey the appropriate

protocols in accessing r, w . A different bag specification is given in Section 4.2 where F is assumed to access r, w appropriately.

3.4 Deducing Properties from the Specification

We deduce one safety and one progress property from the specification given in Section 3.3.

P7. $w = \phi$ unless $|b| < N \vee r = \phi$ in B

In the following proof all properties are in B .

$|r + b + w| \leq N + 1$ constant

, constant formation rule applied to P4

$|r + b + w| \leq N + 1$ stable

, a constant predicate is stable

$w = \phi$ unless $w \neq \phi$

, antireflexivity (see Section 2.1.5)

$w = \phi \wedge |r + b + w| \leq N + 1$ unless $w \neq \phi \wedge |r + b + w| \leq N + 1$

, stable conjunction (Section 2.1.5) applied to the above two

$w = \phi \wedge |r + b| \leq N + 1$ unless $|r + b| \leq N$

, rewriting the lhs and weakening the consequence (Section 2.1.5)

$w = \phi$ unless $|r + b| \leq N$

, substitution axiom applied to the lhs with invariant $|r + b| \leq N + 1$ (see P6)

$w = \phi$ unless $|b| < N \vee r = \phi$

, consequence weakening (Section 2.1.5)

□

P8. If $w = \phi$ stable in F then

$w = \phi \mapsto |b| < N \vee r = \phi$ in $B \parallel F$

In the following proof all properties are in $B \parallel F$ unless otherwise stated.

$w = \phi$ stable in F , given

$w = \phi$ unless $|b| < N \vee r = \phi$ in B , from P7

$w = \phi$ unless $|b| < N \vee r = \phi$ in $B \parallel F$, Corollary 1 of Section 2.2

$|b| > 0 \mapsto w \neq \phi$, from P5.2

$|b| > 0 \wedge w = \phi \mapsto |b| < N \vee r = \phi$, PSP on the above two

$|b| \leq 0 \wedge w = \phi \mapsto |b| < N$, implication rule and using $N > 0$.

$w = \phi \mapsto |b| < N \vee r = \phi$, disjunction on the above two




Figure 2: Concatenation of bags $B1, B2$

3.5 Bag Concatenation

Let $B1$ implement a bag of size $M, M > 0$, with input, output variables, r and v , respectively, and $B2$ implement a bag of size $N, N > 0$, with input, output variables v and w , respectively. We show that $B = B1 \parallel B2$ implements a bag of size $M + N + 1$ with input, output variables r and w , respectively. The arrangement is shown pictorially in Fig. 2. Observe that r and w are different variables (i.e., $B1, B2$ are not connected cyclically) and hence, r cannot be accessed by $B2$ nor can w be accessed by $B1$.

We first rewrite the specification from Section 3.3 for $B1, B2$. In the following, R, V, W are free variables.

There exists $b1$, local to $B1$, such that

- P0. initially $b1 + v = \phi$ in $B1$
- P1. $r \preceq R$ stable in $B1$
- P2. $V \preceq v$ stable in $B1$
- P3. $|b1| \leq M$ invariant in $B1$
- P4. $r + b1 + v$ constant in $B1$
- P5. $\langle \forall G ::$
 - (P5.1) $|b1| < M \mapsto r = \phi$ in $B1 \parallel G$
 - (P5.2) $|b1| > 0 \mapsto v \neq \phi$ in $B1 \parallel G$

There exists $b2$, local to $B2$, such that

- initially $b2 + w = \phi$ in $B2$
- $v \preceq V$ stable in $B2$
- $W \preceq w$ stable in $B2$
- $|b2| \leq N$ invariant in $B2$
- $v + b2 + w$ constant in $B2$
- $\langle \forall H ::$
 - $|b2| < N \mapsto v = \phi$ in $B2 \parallel H$
 - $|b2| > 0 \mapsto w \neq \phi$ in $B2 \parallel H$

The properties of B , to be proven, are (again from Section 3.3)

There exists b , local to B , such that

- P0. initially $b + w = \phi$ in B
- P1. $r \preceq R$ stable in B
- P2. $W \preceq w$ stable in B

P3. $|b| \leq M + N + 1$ invariant in B

P4. $r + b + w$ constant in B

P5. $\langle \forall F \ ::$

(P5.1) $|b| < M + N + 1 \mapsto r = \phi$ in $B \parallel F$

(P5.2) $|b| > 0 \mapsto w \neq \phi$ in $B \parallel F$

\rangle

We start the proof by defining b to be $b1 + v + b2$. To show that b is local to B observe that $b1 + v + b2$ is constant in the environment F since none of the variables— $b1, v$ or $b2$ —can be accessed by F .

Proof of P0: All properties in this proof are in B .

initially $b1 + v = \phi$, from P0 of $B1$

initially $b2 + w = \phi$, from P0 of $B2$

initially $b1 + v + b2 + w = \phi$, from the above two

initially $b + w = \phi$, definition of b □

Proof of P1:

r constant in $B2$, r cannot be accessed by $B2$

$r \preceq R$ stable in $B2$, constant formation; a constant predicate is stable

$r \preceq R$ stable in $B1$, from P1 of $B1$

$r \preceq R$ stable in $B1 \parallel B2$, from Corollary 1 of Section 2.2 □

Proof of P2: Similar to the proof of P1. □

Proof of P3:

$|b1| \leq M$ invariant in $B1 \parallel B2$, from P6 applied to $B1$ (with $F = B2$)

$|v| \leq 1$ invariant in $B1 \parallel B2$, definition of v

$|b2| \leq N$ invariant in $B1 \parallel B2$, from P6 applied to $B2$ (with $F = B1$)

$|b1 + v + b2| \leq M + N + 1$ invariant in $B1 \parallel B2$, from the above three □

Proof of P4: $r + b + w = r + b1 + v + b2 + w$

$r + b1 + v$ constant in $B1$, from P4 of $B1$

$b2 + w$ constant in $B1$, $b2$ and w cannot be accessed by $B1$

$r + b1 + v + b2 + w$ constant in $B1$, constant formation rule

$r + b1 + v + b2 + w$ constant in $B2$, similarly

$r + b1 + v + b2 + w$ constant in $B1 \parallel B2$

, Corollary 2 of Section 2.2

□

Proof of P5: Consider any arbitrary F . Setting G to $B2 \parallel F$ in P5 for $B1$ we get

1. $|b1| < M \mapsto r = \phi$ in $B1 \parallel B2 \parallel F$
2. $|b1| > 0 \mapsto v \neq \phi$ in $B1 \parallel B2 \parallel F$

Similarly, setting H to $B1 \parallel F$ in P5 for $B2$ we get

3. $|b2| < N \mapsto v = \phi$ in $B1 \parallel B2 \parallel F$
4. $|b2| > 0 \mapsto w \neq \phi$ in $B1 \parallel B2 \parallel F$

We will only show P5.1 for B , i.e.

$$|b| < M + N + 1 \mapsto r = \phi \quad \text{in } B \parallel F$$

We use the following fact that is easily seen from the implication and disjunction rules for *leads-to*.

$$\frac{q \mapsto q'}{p \vee q \mapsto p \vee q'}$$

In the following proof all properties are in $B \parallel F$, i.e., $B1 \parallel B2 \parallel F$.

$$\begin{aligned} & |b| < M + N + 1 \\ \Rightarrow & |b1| < M \vee v = \phi \vee |b2| < N && , \text{ since } b = b1 + v + b2 \\ \mapsto & |b1| < M \vee v = \phi && , \text{ above fact about } \mapsto \text{ and (3).} \\ \mapsto & |b1| < M \vee r = \phi && , \text{ above fact about } \mapsto \text{ and P8 for } B1 \text{ (} v = \phi \text{ stable in } B2 \parallel F) \\ \mapsto & r = \phi && , \text{ above fact about } \mapsto \text{ and (1).} \end{aligned}$$

3.6 A Note on the Bag Specification

It is interesting to note that the specification of Section 3.3 does not apply if $N = 0$, i.e., the internal space for storing bag items is empty. In such a case, we would expect program B to move data from r to w whenever $r \neq \phi$ and $w = \phi$. However, the progress condition P5 becomes,

$$\langle \forall F \quad :: \\ \quad |b| < 0 \mapsto r = \phi \quad \text{in } B \parallel F \\ \quad |b| > 0 \mapsto w \neq \phi \quad \text{in } B \parallel F \\ \rangle$$

Since $N = 0$, it follows (from P6 and the definition of b) that $|b| = 0$ is an invariant in $B \parallel F$. Hence, using the substitution axiom, the antecedent of each progress condition is *false*. In UNITY, *false* $\mapsto p$, for any p .

Therefore, for $N = 0$ the specification provides no guarantees on progress. In particular, the specification allows a state $r \neq \phi \wedge w = \phi$ to persist. It may seem that the following strengthening of the progress specification could generalize the specification to $N \geq 0$.

$$\langle \forall F \ :: \\
|b| < N \vee w = \phi \mapsto r = \phi \quad \text{in } B \parallel F \\
|b| > 0 \vee r \neq \phi \mapsto w \neq \phi \quad \text{in } B \parallel F \\
\rangle$$

However since nothing is assumed about F , program F could conceivably change $r \neq \phi \wedge w = \phi$ to $r \neq \phi \wedge w \neq \phi$ (by storing data in w), thus making it impossible for B to implement the first progress condition (while satisfying P4: $r + b + w$ is constant in B). Thus, such a specification cannot be implemented unless we assume something more about the way F behaves, in particular that it never removes data from r nor stores into w . A specification along this line is given in the next section.

4 Alternative Specification of a Bag

4.1 Communication Sequences as Auxiliary Variables

We propose another bag specification in this section. We take as auxiliary variables the bags \hat{r} and \hat{w} , where \hat{r} is the bag of data items written into r (and similarly \hat{w}). Such a bag is typically defined by augmenting the program text appropriately: whenever r is changed \hat{r} is altered by adding the new value of r to it. As shown in an example in Section 2.1.1 our logic provides a direct way of defining \hat{r} , \hat{w} without appealing to the program text. In the following, R, W are free variables (of the same type as data) and X, Y are free variables that are of type bag.

- A0. initially $(\hat{r}, \hat{w}) = (r, w)$ in $B \parallel F$
- A1. $\hat{r} = X \wedge r = R$ unless $\hat{r} = X + r \wedge r \neq R$ in $B \parallel F$
- A2. $\hat{w} = Y \wedge w = W$ unless $\hat{w} = Y + w \wedge w \neq W$ in $B \parallel F$

To understand the relationship between \hat{r} and r and (similarly for \hat{w} and w), note that \hat{r} remains unchanged as long as r is unchanged and \hat{r} may be modified by adding the (new) value of r to it whenever r is changed. Furthermore, in our case, since the values of r will alternate between ϕ and non- ϕ , successive values assumed by r are different. (Note that whenever r is set to ϕ , \hat{r} remains unchanged.)

We deduce two simple facts. (Here “ \subseteq ” is the subbag relation.)

- A3. $r \subseteq \hat{r}$ invariant in $B \parallel F$
- A4. $w \subseteq \hat{w}$ invariant in $B \parallel F$

We prove A3; proof of A4 is similar.

Proof of A3: All properties in the following proof are in $B \parallel F$.

- $\hat{r} = X \wedge r = R$ unless $\hat{r} = X + r \wedge r \neq R$, from A1
- $(\hat{r}, r) = (X, R)$ unless $(\hat{r}, r) \neq (X, R) \wedge r \subseteq \hat{r}$, consequence weakening
- $r \subseteq \hat{r}$ stable , Corollary 1 of Section 2.1.5
- Initially $r \subseteq \hat{r}$ because $r = \hat{r}$. Hence $r \subseteq \hat{r}$ invariant.

4.2 Bag Specification

The following specification is for a bag of size $N, N \geq 0$. The properties R1,R2 are the same as P1,P2. In order to overcome the problems described in Section 3.6, we make some assumptions about the environment, F . Specifically, we require F to obey a similar access protocol to w, r as the ones obeyed by B for access to r, w : Environment F may write into r only if $r = \phi$ and it may change w —to ϕ —only if $w \neq \phi$. The property R3 states that under these conditions on F , the bag \hat{w} is a subbag of $\hat{r} - r$ (property R3.1); these two bags differ by no more than N (property R3.2); if \hat{r}, \hat{w} differ by no more than N then r is or will be set to ϕ (property R3.3); if $|\hat{r}|$ exceeds $|\hat{w}|$ then w is or will be set to non- ϕ (property R3.4).

- R0. initially $w = \phi$ in B
- R1. $r \preceq R$ stable in B
- R2. $W \preceq w$ stable in B
- R3. $\langle \forall F :: B.hypo \text{ in } F \Rightarrow B.conc \text{ in } B \parallel F \rangle$

where

- $B.hypo :: R \preceq r$ stable , $w \preceq W$ stable
- $B.conc ::$
 - R3.1 $\hat{w} \subseteq \hat{r} - r$ invariant
 - R3.2 $|\hat{r} - r - \hat{w}| \leq N$ invariant
 - R3.3 $|\hat{r} - \hat{w}| \leq N \mapsto r = \phi$
 - R3.4 $|\hat{r}| > |\hat{w}| \mapsto w \neq \phi$

We show, in Section 4.3, that under certain conditions the proposed specification is a refinement of the specification in Section 3.3. Now we deduce a few properties from the given specification. Assuming $B.hypo$ in F

- R4. $\hat{r} - \hat{w} + w$ constant in B
- R5. $\hat{r} - \hat{w} - r$ constant in F .

The proof of R5 is similar to that of R4 by interchanging the roles of r, w . We prove R4.

Proof of R4: All properties in the following proof are in B .

$\hat{r} = X \wedge r = R$ unless $\hat{r} = X + r \wedge r \neq R$, union theorem on A1
$r \preceq R$ stable	, from R1
$\hat{r} = X \wedge r = R$ unless $\hat{r} = X + r \wedge r \prec R$, stable conjunction on the above two
$(\hat{r}, r) = (X, R)$ unless $(\hat{r}, r) \neq (X, R) \wedge \hat{r} = X$, consequence weakening
\hat{r} constant	, constant introduction

Similarly, we can show—starting from A2,R2 and applying stable conjunction—that $\hat{w} - w$ is constant in B . Using the constant formation rule, $\hat{r} - (\hat{w} - w)$ is constant in B . Since $w \subseteq \hat{w}$ (from A4) and $\hat{w} \subseteq \hat{r}$ (from R3.1, which can be assumed since $B.hypo$ holds in F), $\hat{r} - (\hat{w} - w) = \hat{r} - \hat{w} + w$, which is then constant in B .

4.3 Proof of Refinement

We show that for $N > 0$, the specification of Section 4.2 is a refinement of the specification of Section 3.3, *provided that $B.hypo$ holds in the environment F* . In the absence of such a requirement on F , the specification of Section 4.2 does not prescribe B 's behavior (that is when B is composed with an arbitrary F) whereas the specification of Section 3.3 prescribes B 's behavior even when it is composed with an arbitrary F . Hence, our proof obligation for the refinement is to show that, given

$B.hypo$ in F , $N > 0$ and the properties R0–R3 of Section 4.2

there exists b , local to B , satisfying properties P0–P5 of Section 3.3. We let

$$b = \hat{r} - \hat{w} - r$$

Observe that b is local to B since b is constant in F (from R5).

Proof of P0: $b + w = \hat{r} - \hat{w} - r + w$

initially $(\hat{r}, \hat{w}) = (r, w)$. Hence initially $b + w = \phi$. □

Proofs of P1, P2: directly from R1, R2. □

Proof of P3:

$|\hat{r} - \hat{w} - r| \leq N$ invariant , from R3.2

$|b| \leq N$ invariant , using the definition of b □

Proof of P4: $r + b + w = \hat{r} - \hat{w} + w$

$\hat{r} - \hat{w} + w$ constant , from R4 □

Proof of P5.1:

$$\begin{aligned}
& |b| < N \wedge r \neq \phi \\
\Rightarrow & |\hat{r} - \hat{w} - r| < N \wedge r \neq \phi && , \text{ using the definition of } b \\
\Rightarrow & |\hat{r} - \hat{w}| \leq N && , r \subseteq \hat{r} \text{ (A3)} \\
\mapsto & r = \phi && , \text{ from R3.3} \\
\text{Hence, } & |b| < N \wedge r \neq \phi \mapsto r = \phi \\
\text{Also, } & |b| < N \wedge r = \phi \mapsto r = \phi && , \text{ implication}
\end{aligned}$$

The result follows using disjunction on the above two. □

The proof of P5.2 is similar to that of P5.1.

5 Specifications of Concurrent Queue and Concurrent Stack

We consider two other concurrent data objects—queue and stack—in this section. Each has variables r, w which are accessed similarly as in the case of the bag. Our interest in these specifications is mainly to show that each of these specifications refines a bag specification, i.e., each implements a bag (of the appropriate size).

5.1 A Specification of a Concurrent Queue

The operation of a concurrent queue is analogous to that of a concurrent bag. The important difference is that in the former case the items are written into w in the same order in which they were read from r . We propose a specification analogous to that of Section 3.3. In the following “ uv ” denotes concatenation of sequences u, v and ϕ denotes the null sequence. A queue of size N , $N > 0$, is defined by a program Q where:

There exists a variable q , local to Q , of type sequence such that

- Q0. initially $q = \phi \wedge w = \phi$ in Q
- Q1. $r \preceq R$ stable in Q
- Q2. $W \preceq w$ stable in Q
- Q3. $|q| \leq N$ invariant in Q
- Q4. rqw constant in Q
- Q5. $\langle \forall F ::$
 - (Q5.1) $|q| < N \mapsto r = \phi$ in $Q \parallel F$
 - (Q5.2) $|q| > 0 \mapsto w \neq \phi$ in $Q \parallel F$

\rangle

5.1.1 Queue Implements Bag

We show that from the specification of Section 5.1 we can deduce the specification in Section 3.3 for an appropriate b . Denote the bag of items in q by $[q]$. Let $b = [q]$.

Proofs of P0,P1,P2 are trivial. Property P3 follows from Q3 and P5 from Q5 by noting that

$$|b| = |[q]| = |q|.$$

To prove P4—that $r + b + w$ constant in B —we observe (below all properties are in B)

$$\begin{array}{ll} rqw \text{ constant} & , \text{ from Q4} \\ [rqw] \text{ constant} & , \text{ constant introduction} \\ [r] + [q] + [w] \text{ constant} & , [rqw] = [r] + [q] + [w] \\ r + b + w \text{ constant} & , [r] = r, [q] = b, [w] = w \end{array}$$

5.2 A Specification of a Concurrent Stack

The operation of a concurrent stack is distinguished from that of a concurrent bag by the requirements that the items read from r be pushed onto a stack (read is permitted only if there is room in the stack) and the item written into w at any point be the top of the stack which is then removed from the stack. A stack is seldom accessed concurrently because speed differences between the producer—process that writes into the stack—and the consumer—that reads from the stack—affects the outcome of the reads.

We propose a specification, analogous to that of Section 4.2, for a program S that implements a stack of size N , $N > 0$. Let \bar{r}, \bar{w} denote the sequences of data items written into r, w respectively. Analogous to the definitions of \hat{r}, \hat{w} by (A0,A1,A2) we define \bar{r}, \bar{w} in $S \parallel F$ (where F is any arbitrary environment of S) by (A0', A1', A2') below. As in Section 5.1, concatenations of sequences are shown by juxtapositions and ϕ denotes the null sequence.

$$\begin{array}{l} \text{(A0')} \quad \text{initially } (\bar{r}, \bar{w}) = (r, w) \quad \text{in } S \parallel F \\ \text{(A1')} \quad \bar{r} = X \wedge r = R \quad \text{unless } \bar{r} = Xr \wedge r \neq R \quad \text{in } S \parallel F \\ \text{(A2')} \quad \bar{w} = Y \wedge w = W \quad \text{unless } \bar{w} = Yw \wedge w \neq W \quad \text{in } S \parallel F \end{array}$$

Now, we define $X N Y$ —a binary relation between sequences X, Y expressing that Y is the complete output sequence of a stack of size N given X as the input sequence. For $N \geq 0$, it is the strongest relation satisfying

- $\phi N \phi$
- $X N X', Y (N + 1) Y' \Rightarrow aXY (N + 1) X'aY'$

where a is any arbitrary data item. The second rule states that given an input string aXY to a stack of size $N + 1$, the item a appears in the output at some point. Prior to output of a , item a is at the bottom of the stack and hence, the stack behaves as if its size is N in converting some portion of the input, X , to X' . Following the output of a , the stack is empty and hence the remaining input sequence Y is converted to Y' using a stack of size $N + 1$.

Several interesting properties—using induction on the length of X —can be proven about this relation. In particular,

$$X N Y \Rightarrow [X] = [Y]$$

where $[X]$ is the bag of items in X .

Now we specify program S that implements a concurrent stack of size N , $N > 0$, with shared variables r, w as before. Note that S3.1 is the only property that differs substantially from the corresponding property of the bag in Section 4.2.

- S0. initially $w = \phi$ in S
- S1. $r \preceq R$ stable in S
- S2. $W \preceq w$ stable in S
- S3. $\langle \forall F :: B.hypo \text{ in } S \Rightarrow B.conc \text{ in } S \parallel F \rangle$

where

- $B.hypo :: R \preceq r \text{ stable}, w \preceq W \text{ stable}$
- $B.conc :: \{ \text{In S3.1, } \bar{r} - r \text{ is the prefix of } \bar{r} \text{ excluding } r \text{ (if } r = \phi, \bar{r} - r = \bar{r} \text{)}. \text{ Also, } \sqsubseteq \text{ is the prefix relation over sequences} \}$
- S3.1 $\langle \exists \bar{z} :: (\bar{r} - r) N \bar{z} \wedge \bar{w} \sqsubseteq \bar{z} \rangle$ invariant
- S3.2 $|\bar{r}| \leq |r| + |\bar{w}| + N$ invariant
- S3.3 $|\bar{r}| \leq |\bar{w}| + N \mapsto r = \phi$
- S3.4 $|\bar{r}| > |\bar{w}| \mapsto w \neq \phi$

5.2.1 Stack Implements a Bag

We show that from the specification of a concurrent stack we can deduce the bag specification in Section 4.2 for $N > 0$. Note that $\hat{r} = [r]$, $\hat{w} = [\bar{w}]$. It follows that $\hat{r} - r = [\hat{r} - r]$. Proofs of properties R0–R3 are entirely straightforward, except for R3.1: $\hat{w} \subseteq \hat{r} - r$.

$$\begin{array}{lll}
(\bar{r} - r) N \bar{z} & \Rightarrow [\bar{r} - r] = [\bar{z}] & , \text{ property of the binary relation stated earlier} \\
\bar{w} \sqsubseteq \bar{z} & \Rightarrow [\bar{w}] \subseteq [\bar{z}] & , \text{ fact about building a bag from a sequence} \\
(\bar{r} - r) N \bar{z} \wedge \bar{w} \sqsubseteq \bar{z} & \Rightarrow [\bar{w}] \subseteq [\bar{r} - r] & , \text{ from the above two}
\end{array}$$

$\langle \exists \bar{z} :: (\bar{r} - r) N \bar{z} \wedge \bar{w} \sqsubseteq \bar{z} \rangle \Rightarrow \hat{w} \subseteq \hat{r} - r$, from the above using $[\bar{w}] = \hat{w}$ and $[\bar{r} - r] = \hat{r} - r$.

We leave it as an exercise for the reader to show that a queue of size 1 implements a stack of size 1, and vice versa.

6 Refinement of the Bag Specification

We refine the specification of Sec. 4.2 as a step toward implementing a bag. It can be shown from that specification that concatenation of two bags of size M, N , where M, N are non-negative, results in a bag of size $M + N + 1$; the proof is similar in structure to the proof in Sec. 3.5; also, a similar proof for a queue appears in [7]. Hence a bag of size N , $N > 0$, can be implemented by concatenating $(N + 1)$ bags of size 0 each. Therefore we restrict our attention to bags of size 0 and propose a refinement. If a bag has size 0, the only possible action is to move data items from r to w directly. The effect of this action is to keep $r + w$ constant (property T3), and such an action is guaranteed to take place provided $r \neq \phi \wedge w = \phi$ resulting in $w \neq \phi$ (property T4).

- T0. initially $w = \phi$ in B
- T1. $r \preceq R$ stable in B
- T2. $W \preceq w$ stable in B
- T3. $r + w$ constant in B
- T4. $r \neq \phi$ ensures $w \neq \phi$ in B

6.1 Proof of the Refinement

We show that the properties R0–R3 of Sec. 4.2, for $N = 0$, can be deduced from T0–T4. Proofs of R0,R1,R2 are immediate from T0,T1,T2. Property R3 is of the form,

$$\langle \forall F :: B.hypo \text{ in } F \Rightarrow B.conc \text{ in } B \parallel F \rangle$$

We first show that

$$\text{T5: } \langle \forall F :: B.hypo \text{ in } F \wedge (T0 \wedge T3) \Rightarrow \hat{r} = \hat{w} + r \text{ invariant in } B \parallel F \rangle$$

Proof:

$$\begin{array}{ll} \hat{r} - \hat{w} + w \text{ constant in } B & , \text{ from R4} \\ r + w \text{ constant in } B & , \text{ from T3} \\ \hat{r} - \hat{w} + w - (r + w) \text{ constant in } B & , \text{ constant formation} \\ \hat{r} - \hat{w} - r \text{ constant in } B & , \text{ simplifying the above expression} \\ \hat{r} - \hat{w} - r \text{ constant in } F & , \text{ from R5 assuming } B.hypo \text{ in } F \end{array}$$

$\hat{r} - \hat{w} - r$ constant in $B \parallel F$, Corollary 2 of Sec. 2.2
 Initially $\hat{r} - \hat{w} - r = \phi$ in $B \parallel F$, from $\hat{r} = r$ and $\hat{w} = w = \phi$
 Hence $\hat{r} - \hat{w} - r = \phi$ invariant in $B \parallel F$, i.e., $\hat{r} = \hat{w} + r$ invariant in $B \parallel F$. \square

Next we prove the four properties in $B.conc$, from $B.hypo$ in F , T1, T2, and $\hat{r} = \hat{w} + r$ invariant in $B \parallel F$.

Proof of R3: Consider an arbitrary F in which $B.hypo$ holds.

R3.1. Proof is immediate given T5. \square

R3.2. Proof is immediate given T5 and that $N = 0$ \square

R3.3. We have to show that $|\hat{r}| \leq |\hat{w}| \mapsto r = \phi$ in $B \parallel F$.

In the following all properties are in $B \parallel F$.

$|\hat{r}| \leq |\hat{w}| \Rightarrow r = \phi$, from T5

$|\hat{r}| \leq |\hat{w}| \mapsto r = \phi$, using implication rule on the above \square

R3.4. We have to show that $|\hat{r}| > |\hat{w}| \mapsto w \neq \phi$ in $B \parallel F$

$r = R$ unless $r \neq R$ in F , antireflexivity of *unless*

$r = R$ unless $r \neq R \wedge r \neq \phi$ in F , stable conjunction with $R \preceq r$ stable (from $B.hypo$)

$r \neq \phi$ stable in F , Corollary 1 of Sec. 2.1.5

$r \neq \phi$ ensures $w \neq \phi$ in B , from T4

$r \neq \phi$ ensures $w \neq \phi$ in $B \parallel F$, using Corollary 4 of Sec. 2.2 on the above two

$r \neq \phi \mapsto w \neq \phi$ in $B \parallel F$, definition of \mapsto

$|\hat{r}| > |\hat{w}| \mapsto w \neq \phi$ in $B \parallel F$, from T5 $|\hat{r}| > |\hat{w}| \equiv r \neq \phi$

7 An Implementation

The specification of Section 6 can be implemented by a program whose only statement moves data from r to w provided $w = \phi$ (if $r = \phi$, the movement has no effect):

$$r, w := \phi, r \quad \text{if } w = \phi$$

The proof that this fragment has the properties T1–T4 is immediate from the definition of *unless* and *ensures*. The initial condition of this program is $w = \phi$, and hence (T0) is established.

An implementation for a bag of size N , $N > 0$, uses the union of $N + 1$ such statements: One statement each for moving data from a location to an adjacent location (closer to w) provided the latter is ϕ . We show how this program may be expressed in the UNITY programming notation.

Rename the variables r and w to be $b[0]$ and $b[N + 1]$, respectively. The internal bag words are $b[1]$ through $b[N]$. In the following program we write $\langle \parallel i : 0 < i \leq N + 1 \quad :: \quad t(i) \rangle$ as a shorthand for $t(1) \parallel t(2) \dots \parallel t(N + 1)$,

where $t(1)$, for instance, is obtained by replacing every occurrence of i by 1 in $t(i)$. The program specifies the initial values of $b[1]$ through $b[N + 1]$ to be ϕ (in the part followed by **initially**). The statements of the program are given after **assign**; the generic statement shown moves $b[i - 1]$ to $b[i]$ provided the latter is ϕ .

Program *bag* {of size N , $N > 0$ }

initially $\langle \parallel i : 0 < i \leq N + 1 \ :: b[i] = \phi \rangle$

assign $\langle \parallel i : 0 < i \leq N + 1 \ :: b[i - 1], b[i] := \phi, b[i - 1] \quad \text{if } b[i] = \phi \rangle$

end {*bag*}

8 Evaluation of the Proposed Methodology

Equational notation, as in [2], is generally preferred for specifications because it enables a programmer to separate two concerns: deducing facts about an object independent of its implementation, and implementing the given equations to minimize time or storage. As we have argued, the objects that are concurrently accessed are, as yet, not amenable to equational specifications. An important research problem is to determine the conditions under which an object can be specified equationally, and yet accessed concurrently.

Our specification of the bag illustrates how some of the requirements described in the introduction are met. The specification is concise even though it makes concurrent access explicit. The notation admits of efficient manipulation—see, for instance, the lengths of the refinement proofs in Sections 4, 5 and 6. We believe that one of the most important applications of formal specifications is reasoning about an object; a cumbersome notation—such as a programming notation—defeats the main purpose of specification.

It may seem from our specification that we have artificially simplified the problem by introducing shared variables— r, w in case of bag, queue, and stack—through which the object communicates with this environment. How about multiple processes accessing an object through remote procedure calls, for instance? First, remote procedure call by a process is equivalent to the process storing the procedure parameters in certain (shared) variables which the object can later access; in this sense, we have not introduced any artificial simplicity. Second, accesses by multiple processes—for instance, multiple producers writing into shared variables $r.j$, $0 \leq j < L$ —can be treated within our theory as follows. Imagine that a *sequencer* process copies some $r.j$ into r whenever $r = \phi$; the different $r.j$'s are chosen fairly during the course of the program execution. The index j could be part of the data written into r , if, for instance, the object needs to respond to a request in $r.j$ by writing into $w.j$. (Similarly, outputs for different consumers may be sequenced.) The entire concurrent object may be viewed as a union of an input sequencer, data manager (the program that actually implements the data structures) and the output sequencer, and each of these may be specified individually. We have shown how to specify the data manager; specifications of input/output sequencers are entirely straightforward. A direct specification of a queue with multiple input/output variables appears in [8].

9 Acknowledgment

It was a suggestion from Leslie Lamport which led me to study this problem and to understand how auxiliary variables can be defined and effectively used. I am indebted to Ambuj Singh for suggesting simplifications for my original specification; his work clearly reveals the importance of formal reasoning as I was unable to justify his simplifications using intuitive arguments alone. Comments from the participants of the 9th International Summer School in Marktoberdorf, W. Germany, were most helpful; particular thanks go to Jan L. A. van de Snepscheut. Thorough and critical readings of different versions of the manuscript by Mani Chandy, Edgar Knapp and J. R. Rao have improved the presentation.

10 References

1. K. Mani Chandy and Jayadev Misra, *Parallel Program Design: A Foundation*, Addison-Wesley, Reading, MA, 1988.
2. John Guttag, "Abstract Data Types and the Development of Data Structures," *C.ACM*, 20:6, pp. 396–404, June 1977.
3. B. Hailpern, "Verifying Concurrent Processes Using Temporal Logic," Springer-Verlag, *Lecture Notes in Computer Science 129*, 1982.
4. *Specification Case Studies*, Ian Hayes (ed.), Prentice-Hall International, Englewood Cliffs, N.J., 1987.
5. L. Lamport, "A Simple Approach to Specifying Concurrent Systems," *Digital Systems Research Center Report 15*, December 1986, 130 Lytton Avenue, Palo Alto, CA 94301.
6. J. Misra, "Monotonicity, Stability and Constants," *Notes on Unity 10-89*, University of Texas at Austin, 1989.
7. J. Misra, "Specifications of Concurrently Accessed Data," *Proc. International Conference on Mathematics of Program Construction, Groningen University, The Netherlands, June 26–30, 1989, LNCS 375*, Springer-Verlag.
8. Ambuj Singh, "Ranking in Distributed Systems," (see Sec. 3.4, in particular), Ph.D. thesis, The University of Texas at Austin, 1989.