

A CRITERION FOR ATOMICITY*

James H. Anderson and Mohamed G. Gouda

Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712-1188

TR-90-24

July 1990

ABSTRACT

Most proof methods for reasoning about concurrent programs are based upon the *interleaving semantics* of concurrent computation: a concurrent program is executed in a stepwise fashion, with only one enabled action being executed at each step. Interleaving semantics, in effect, requires that a concurrent program be executed as a nondeterministic sequential program. This is clearly an abstraction of the way in which concurrent programs are actually executed. To ensure that this is a reasonable abstraction, interleaving semantics should only be used to reason about programs with “simple” actions; we call such programs “atomic.” In this paper, we formally characterize the class of atomic programs. We adopt the criterion that a program is *atomic* if it can be implemented in a wait-free, serializable manner by a primitive program. A program is *primitive* if each of its actions has at most one occurrence of a shared bit, and each shared bit is read by at most one process and written by at most one process. It follows from our results that the traditionally accepted atomicity criterion, which allows each action to have at most one occurrence of a shared variable, can be relaxed, allowing programs to have more powerful actions. For example, according to our criterion, an action can read any finite number of shared variables, provided it writes no shared variable.

Keywords: Atomicity, atomic registers, composite registers, history, interleaving semantics, leader election, models of concurrency, shared variable, snapshot, waiting.

CR Categories: D.4.1, D.4.2, D.4.4, F.1.1, F.1.2.

* This work was supported in part by Office of Naval Research Contract N00014-89-J-1913.

1 Introduction

Most proof methods that have been proposed for reasoning about concurrent programs are based upon the assumption that only one action of a concurrent program is executed at a time [14, 18, 21, 24]. This model of program execution is often referred to as *interleaving semantics* since the actions executed by one process are interleaved with those of other processes. Interleaving semantics is desirable because it allows us to view a concurrent program as being equivalent to a nondeterministic sequential program. Therefore, techniques that have been developed for reasoning about sequential programs (such as using invariants to prove safety properties and well-founded rankings to prove progress properties) can be applied to concurrent programs as well.

Interleaving semantics is clearly an abstraction of the way in which concurrent programs are actually executed. In order to fully exploit the power of a parallel machine, concurrent programs are usually implemented in a manner that allows the executions of several actions to overlap in time. To ensure that interleaving semantics is a suitable abstraction, such an implementation is typically expected to satisfy the requirement that every parallel execution of two or more actions be “equivalent” to some interleaved one. For example, if a read and a write of a variable overlap in time, then the read either obtains the value before the write or the value after the write.

The task of implementing a concurrent program so that any parallel execution is equivalent to some interleaved one may be difficult, unless we restrict our attention to programs with sufficiently simple actions. For example, consider a concurrent program consisting of processes P_0, \dots, P_{99} and variables X_0, \dots, X_{99} ; each process P_i , where $0 \leq i < 100$, consists of a single action $X_i := (\text{SUM } j : 0 \leq j < 100 : X_j)$. If this program is to be executed in accordance with interleaving semantics, then a high degree of synchronization and coordination among processes will be required. We contend that such a program will not be “easy” to implement on a parallel machine.

To ensure that programs are easily implementable, interleaving semantics should not be used to reason about all programs, but only those programs with sufficiently simple actions; we call such programs “atomic.” The purpose of this paper is to give a criterion for determining those programs that should be considered atomic.

If ease of implementation were our only concern, then we could stop here and define the class of atomic programs to include only those programs with the very simplest of actions. However, ease of implementation is not our only concern. We would also like to ease the burden on the programmer as much as possible by allowing a rich set of actions. Thus, we are faced with a dilemma: For ease of implementation, we should be as strict as possible, but for ease of programming, we should be as lax as possible. How can we possibly satisfy both of these seemingly conflicting requirements?

The roots of this question can be traced back to the seminal work of Owicki and Gries [24], who were perhaps the first to define a criterion for atomicity. In order to state their atomicity criterion, it is necessary to first distinguish between two classes of program variables, shared and private. A variable is *shared* if it appears among the actions of more than one process; otherwise, it is *private*. Owicki and Gries adopted the criterion that each action of a concurrent program has at most one occurrence of at most one shared variable. For example, if a is private and X and Y are shared,¹ then the action $X := a + 15$ would be allowed by their criterion, but the actions $X := Y + 2$ and

¹ We will generally follow the convention of using upper-case letters for shared variables, and lower-case letters for private variables.

$Y := Y + 1$ would not. The underlying assumption here is that it should be possible to read or write any shared variable in a single action. Owicki and Gries attribute the “observation that the memory reference must have ‘reasonable’ properties” to John Reynolds [24]. For this reason, the Owicki-Gries atomicity criterion is sometimes known as Reynolds’ Rule [17]. This atomicity criterion has gained widespread acceptance since its introduction.

Owicki and Gries, in effect, resolve the dilemma noted above “by definition” because they do not defend their atomicity criterion on any formal basis. We are thus led to question their criterion on two grounds. First, it is not clear whether this is a viable criterion from an implementation standpoint. In fact, its “reasonableness” in this regard probably qualifies as a folk theorem [17]. Second, we question whether this criterion is unnecessarily restrictive. That is, it might be possible to relax the Owicki-Gries atomicity criterion somewhat without making programs any harder to implement.

In this paper, we remedy these two shortcomings by providing a formal framework for defining the class of atomic programs. This formal framework is based upon three concepts: primitive actions, wait-freedom, and serializability. We introduce these concepts now by returning to the central question posed above: What kinds of actions should an atomicity criterion allow, given the conflicting requirements of ease of implementation and ease of programming?

Clearly, an atomicity criterion should allow one process to communicate with another, i.e., it is unreasonable to disallow all actions that access shared variables. We therefore choose to allow some actions that access shared variables, but to ensure that programs are easily implementable, we will begin by insisting that each such action be as restrictive as possible. This gives rise to the notion of a “primitive” action: an action is *primitive* if it has at most one occurrence of a shared bit, and each shared bit is read by at most one process and written by at most one process. For example, let a and b be private bits and let X and Y be shared bits that are read by at most one process and written by at most one process. Then, the action $a, b := X, \neg b$ would be considered primitive, but the actions $a, b := X, \neg X$ and $X := Y$ would not. Our starting point, then, is the assumption that each primitive action is reasonable enough to implement.

Given this assumption, it seems reasonable to allow certain nonprimitive actions as well. In particular, we also choose to allow any action that can be implemented in terms of primitive actions without a high degree of synchronization or coordination among processes. We illustrate this by considering two examples. First of all, consider the first nonprimitive action given above, i.e., $a, b := X, \neg X$. This action can be implemented in terms of primitive actions by introducing a new private bit c as follows.

$$c := X; \quad a, b := c, \neg c$$

Such an implementation seems sufficiently simple, so we conclude that our atomicity criterion should also allow an action such as $a, b := X, \neg X$.

Next, consider the other nonprimitive action mentioned above, i.e., $X := Y$. It turns out that this action can be quite difficult to implement in terms of primitive actions. In particular, consider the following concurrent program, where initially $X \neq Y$.

$$\text{process } P : X := Y \qquad \text{process } Q : Y := X$$

According to interleaving semantics, either process P executes its action first or process Q does. In either case, the program terminates in a state satisfying the assertion $X = Y$. However, as shown

in Section 5, if we implement each action $X := Y$ and $Y := X$ in terms of primitive actions, then at least one of the two processes P and Q must busy-wait in order to enforce the necessary mutual exclusion. A process that busy-waits may execute an infinite number of primitive actions, and hence may never terminate. Thus, it is impossible to implement the actions of this program in terms of primitive actions so that the program always terminates in a state satisfying $X = Y$. We therefore conclude that it is unreasonable to allow an action such as $X := Y$.

With this motivation in mind, our atomicity criterion can be described as follows. We define a program to be *atomic* if it can be implemented in a wait-free, serializable manner by a primitive program. A program is *primitive* iff each of its actions is primitive. Our notion of serializability is reflected by the implementation of the action $a, b := X, \neg X$ given above. In particular, this notion requires that whenever two or more implemented actions are executed concurrently (i.e., the primitive actions that implement them are interleaved), there exists an equivalent sequential execution.

Any program allowed by our atomicity criterion is easily implementable because it can be “reduced” to one with only primitive actions. It turns out that our atomicity criterion includes all programs allowed by the Owicki-Gries atomicity criterion.² This confirms their choice of actions from an implementation standpoint. However, as we shall see shortly, our atomicity criterion allows some rather complicated actions that are not allowed by the Owicki-Gries criterion. For example, our criterion allows an action to read any finite number of shared variables, provided it writes no shared variable. Thus, our results establish that the traditional Owicki-Gries atomicity criterion is too restrictive.

The rest of this paper is organized as follows. In Section 2, we present our model of concurrent programs, and in Section 3, we present the proposed atomicity criterion mentioned above. The task of checking whether a program is atomic is nontrivial if we appeal directly to this criterion. Therefore, we present a method that allows this task to be performed systematically. We outline our approach to justifying this method at the end of Section 3. The actual justification is given in Sections 4 through 8. Concluding remarks appear in Section 9.

2 Concurrent Programs

A *concurrent program* consists of a set of processes and a set of variables. A *process* is a sequential program specified using Dijkstra’s guarded commands [10]. Each *variable* is either unsubscripted or subscripted. Unsubscripted variables are defined using integer, boolean, and subrange (e.g., $0..K$) types. Subscripted variables are defined using an array type; the elements of an array are of type integer, boolean, or subrange. Multi-dimensional arrays and arrays with an infinite number of elements are allowed.

Each variable of a concurrent program is either private or shared. A *private variable* is defined only within the scope of a single process, whereas a *shared variable* is defined globally and may be accessed by more than one process. We assume that the elements of an array are either all shared

²At least, for the kinds of programs considered here. We require processes to be specified using ordinary sequential programming statements, whereas Owicki and Gries also allow the possibility of an *await* statement. Such a statement obviously cannot be implemented in a wait-free fashion, and hence should not be allowed.

or all private to the same process. We also make the simplifying assumption that only expressions over private variables are used as array subscripts.

In the remainder of this paper, the term “variable” is meant to refer to either an unsubscripted variable or an element of an array — we do not refer to an entire array as being a single variable.

The processes of a concurrent program are allowed to access shared variables only by executing assignment statements. An assignment statement has the form

$$x_1, \dots, x_N := E_1, \dots, E_N$$

where each x_i is a distinct variable and each E_j is an expression over a finite number of variables. This statement is executed as a single action as follows: first, the expressions E_1, \dots, E_N are evaluated along with any array subscripts appearing in x_1, \dots, x_N ; then, the value of each expression E_i is assigned to the corresponding variable x_i . We say that the variables appearing in the expressions E_1, \dots, E_N and the array subscripts appearing in x_1, \dots, x_N are *read* by the assignment statement, and the variables x_1, \dots, x_N are *written* by the assignment statement. For example, consider the assignment statement $X[i], i := Y[j], j + 1$. This statement reads the variables i, j , and $Y[j]$, and writes the variables i and $X[i]$.

Associated with each shared variable is a set of processes that are allowed to *read* the variable, and a set of processes that are allowed to *write* the variable. These two sets of processes must be specified as part of the definition of every shared variable. An assignment statement of a process may read (write) a shared variable only if that process has read (write) access to that variable. (Presumably, a runtime error would occur if this rule is ever violated.)

3 Atomicity Criterion

In this section, we define two important classes of programs, “primitive” and “atomic.” Informally, a program is primitive if its processes communicate only by means of single-reader, single-writer, shared bits. More precisely, a program is *primitive* iff it satisfies the following six constraints.

- (0) Each shared variable is a single bit.
- (1) Each shared variable is read by at most one process.
- (2) Each assignment statement reads at most one shared variable and does so at most once.
- (3) Each shared variable is written by at most one process.
- (4) Each assignment statement writes at most one shared variable.
- (5) An assignment statement does not both read a shared variable and write a shared variable.

Note that these constraints can be checked syntactically. In particular, constraint (0) can be checked by simply noting the type of each shared variable. Constraints (1) and (3) can be checked by counting the number of processes that are allowed to read or write each shared variable. Constraints (2), (4), and (5) can be checked by counting the number of occurrences of shared variables in each assignment statement. This can be done on a syntactic basis, given our assumptions concerning

variable declarations. (For example, consider the statement $a, b := X[i], X[j]$, where a and b are private. By assumption, either all elements of the array X are shared, or all are private to the same process. If each element of X is shared, then there are two occurrences of shared variables in this statement, namely $X[i]$ and $X[j]$; if, on the other hand, each element of X is private, then there are no occurrences of shared variables in this statement.)

As stated previously, we assume that any primitive program is easily implementable. Our definition of the more general class of atomic programs reflects this assumption. Simply stated, this definition is as follows.

Atomicity Criterion: A program is *atomic* iff it can be implemented by a primitive program.

The notion of “implements” referred to here is formally defined in Section 4. Informally, we implement an atomic program by replacing each assignment statement that fails to satisfy constraints (0) through (5) by a program fragment whose statements satisfy (0) through (5). Different program fragments in different processes may be executed concurrently (i.e., their statements may be interleaved); each such execution, however, is required to “resemble” one in which the program fragments are executed in sequence. The program fragments are restricted to be wait-free, i.e., busy-waiting loops that can iterate an infinite number of times are not allowed. This restriction guarantees that a process executes each of its implemented assignment statements in a finite amount of time, regardless of the activities of other processes.

In general, the task of proving that a given program satisfies our proposed atomicity criterion is nontrivial. A proof must establish that whenever two or more of the above-mentioned program fragments in different processes are executed concurrently, there exists an equivalent execution in which they are executed in sequence. As we shall see later, the number of different cases that must be taken into account in order to establish this proof obligation can be considerable.

The main contribution of this paper is to give a simple method for checking whether a given program satisfies the proposed atomicity criterion. In many cases, this method amounts to a syntactic check of the program text. (For instance, see the example below.) In particular, we prove that a program is atomic if its assignment statements satisfy the following three conditions.

- *Exact-Writing:* The shared variables can be partitioned into classes, where each class either consists of a single variable, or a number of unsubscripted variables, such that the following condition holds: if an assignment statement writes a variable from some class, then it writes every variable in that class and no variable from any other class.
- *Single-Phase:* If an assignment statement of one process reads a variable that is written by another process, then that assignment statement does not write a shared variable.
- *Finite-Reading:* If an assignment statement reads an element of an infinite shared array, then it reads no other shared variable.

We call these three conditions *atomicity constraints*. The Exact-Writing and Finite-Reading constraints can be checked on a syntactic basis by simply checking the shared variables that are written and read, respectively, by each assignment statement. By contrast, it may be necessary to take the

semantics of a program into account in order to check whether Single-Phase is satisfied. For example, consider the assignment statement $X[i] := Y[j]$ of some process p . To determine whether this statement violates Single-Phase, it is necessary to determine whether $Y[j]$ is written by a process other than p — but, this may depend on the value of j . Note that we are, in effect, viewing each statement such as $X[i] := Y[j]$ as being a separate statement for each of the possible values of i and j .

The Exact-Writing constraint specifies the conditions under which an assignment statement may write several shared variables. Intuitively, this constraint implies that we can view the set of shared variables written by a given assignment statement as a single “compound” variable. (For the sake of simplicity, we have made Exact-Writing more restrictive than is necessary. In particular, we allow a class to consist of more than one shared variable only if each variable in the class is unsubscripted. In [4], a more permissive version of Exact-Writing is considered in which a class may consist of several subscripted variables. This version of Exact-Writing is considered in Section 9.) The Single-Phase constraint specifies the conditions under which an assignment statement may both read and write shared variables. As we shall see in Section 6, this constraint implies that any assignment statement can easily be implemented in terms of assignment statements that either read shared variables or write shared variables; hence the name “Single-Phase.” The Finite-Reading constraint specifies the conditions under which an assignment statement may read several shared variables. Note that this is a very permissive constraint, as it only restricts those programs that have infinite shared arrays.

We illustrate the atomicity constraints by the following example.

Example: Consider the concurrent program in Figure 1. This program consists of a *controller* process and N *device* processes that communicate by means of five shared variables V , W , X , Y , and Z . Variables V , X , and Y are written by the *controller*, and variables W and Z are written by the *devices*. Variables V and X are read by the *devices*, variable W is read by the *controller*, and variables Y and Z are read by all $N + 1$ processes.

The five variables V , W , X , Y , and Z are used as follows. V stores a “request” as produced by the *controller* and X identifies the *device* for which the request is intended. The particular *device*’s “reply” to the given request is stored in W . A request is “pending” iff $Y \neq Z$; Y is updated by the *controller* and Z by the *devices*. The protocol followed by these processes is simple. The *controller* repeatedly consumes replies and produces requests, and each *device* repeatedly consumes requests and produces replies.

To verify that Exact-Writing is satisfied, we partition the shared variables into classes as follows: the tuple $\{V, X, Y\}$ is a class, and the pair $\{W, Z\}$ is a class. Observe that the only assignment statements that write shared variables are

$$V, X, Y := request, dev, \neg Y$$

and

$$W, Z := reply, z .$$

Thus, because V , X , and Y are always written together, and because W and Z are always written together, Exact-Writing is satisfied. Note that the only assignment statement that both reads and

```

shared var V, W : integer;
        X : 1..N;
        Y, Z : boolean

initialization (W = default value) ∧ (Y = Z)

process controller
private var request, reply : integer;
        dev : 1..N;
        pending : boolean
begin
do true → reply, pending := W, (Y ≠ Z);
    if ¬pending → consume reply; produce next request, dev;
        V, X, Y := request, dev, ¬Y
    [] pending → skip
    fi
od
end

process device(j : 1..N)
private var request, reply : integer;
        dev : 1..N;
        pending, z : boolean
begin
do true → request, dev, pending := V, X, (Y ≠ Z);
    if pending ∧ dev = j → consume request; produce reply;
        z := ¬Z;
        W, Z := reply, z
    [] ¬pending ∨ dev ≠ j → skip
    fi
od
end

```

Figure 1: An example program.

writes shared variables is the statement

$$V, X, Y := \text{request}, \text{dev}, \neg Y$$

of the *controller*. Because variable Y is written only by the *controller*, this implies that Single-Phase is satisfied. Finite-Reading is trivially satisfied because there are no infinite shared arrays. Notice that this program violates the traditional atomicity criterion of Owicki and Gries. \square

The remaining sections are dedicated to justifying the proposed atomicity constraints: Exact-Writing, Single-Phase, and Finite-Reading. In particular, we show that any program satisfying these

atomicity constraints is atomic. We also show that Exact-Writing and Single-Phase are necessary in the following sense: for each of these constraints, there exists a nonatomic program violating that constraint, but satisfying the other two constraints. This shows that if we remove either of these two constraints from our set of constraints, then we can no longer make the claim that every program satisfying our set of constraints is atomic. As yet, we do not know whether Finite-Reading is necessary in this sense. However, this constraint does not appear to be a severe limitation, as it only restricts those programs that have infinite shared arrays.

To establish the necessity of Exact-Writing and Single-Phase, we consider in Section 5 the problem of two-process leader election. We first show that there is no primitive program that solves this problem. Then, we exhibit two programs that solve this problem, one that violates Exact-Writing and satisfies Single-Phase and Finite-Reading, and another that violates Single-Phase and satisfies Exact-Writing and Finite-Reading. Neither of these two programs can be implemented by a primitive program, because we know that such a program does not exist. Thus, both programs are nonatomic. This establishes the following theorem.

Theorem 1: There exists a nonatomic program violating Exact-Writing (or Single-Phase), but satisfying the other two constraints. \square

To establish the sufficiency of our atomicity constraints, we consider in Section 6 a shared data object called a *composite register* [2, 3]. A composite register is an array-like variable consisting of a number of components. An operation of the register either writes a value to a single component, or reads the values of all of the components. As shown in Section 6, any program satisfying our proposed atomicity constraints can be implemented by a program with processes that communicate only by reading and writing composite registers. Thus, we can establish the sufficiency of our atomicity constraints by showing that a composite register can be implemented in terms of single-reader, single-writer, shared bits. Given the structure of a composite register, this proof obligation can be broken into several steps. The register constructions for these steps are described in Section 6; one of the constructions is presented in detail in Section 8. These constructions establish the following theorem.

Theorem 2: Any program satisfying Exact-Writing, Single-Phase, and Finite-Reading is atomic. \square

4 Implementing Atomic Programs

As stated in Section 3, a program is atomic iff it can be implemented by a primitive program. In this section, we formally define what it means to implement one program by another.

Terminology: If program P is implemented by program Q , then we refer to P as the *implemented program*, and Q as the *implementation*. \square

Definition: An *implementation* is obtained by replacing each shared variable of the implemented

program by a set of variables, and each assignment statement of the implemented program by a program fragment. In principle, such a program fragment can be expressed as follows:

$$S_0; \text{ do } B_1 \rightarrow S_1 \parallel \cdots \parallel B_N \rightarrow S_N \text{ od}$$

where each B_i is an expression over private variables, and each S_j is an assignment statement. \square

As we shall see later, an implementation is *correct* iff it is wait-free, sequentially correct, and serializable. Informally, these conditions are defined as follows.

- *Wait-Freedom*: Each program fragment is free of busy-waiting loops that can potentially iterate an infinite number of times.
- *Sequential Correctness*: When executed in isolation, each program fragment has the same “effect” as the assignment statement that it replaces.
- *Serializability*: Whenever two or more program fragments in different processes are executed concurrently (i.e., their statements are interleaved), there exists an “equivalent” execution in which the program fragments are executed in sequence.

In this section, we formally define these conditions. We motivate our formal treatment by considering the following example. As this example shows, it is usually straightforward to define an implementation so that wait-freedom and sequential correctness hold. The principal difficulty lies in satisfying serializability.

Example: Consider a concurrent program with a process that includes the assignment statement $X := 4$, where X is a single-reader, single-writer, four-bit, shared variable. Note that this statement violates constraint (0) in the definition of primitive programs given in Section 3. To implement the given program by a primitive program, it is necessary to replace X by a set of single-bit variables. One way to do this is to partition variable X into four separate single-bit variables X_0 , X_1 , X_2 , and X_3 . Then, we can replace the statement $X := 4$ by the following program fragment.

```

X0 := 0;
X1 := 0;
X2 := 1;
X3 := 0

```

Each other assignment statement of the implemented program that reads (or writes) X would similarly be replaced by a program fragment that reads (or writes) each X_i .

The above program fragment for $X := 4$ is loop-free. Moreover, executing this program fragment in isolation clearly has the effect of assigning the value 4 to variable X . Thus, this implementation is both wait-free and sequentially correct. However, a read of X may obtain an incorrect value if it is executed concurrently with a write to X . Suppose, for example, that the program fragment for $X := 4$ is executed when $X = 15$ (i.e., when $X_0 = X_1 = X_2 = X_3 = 1$). If a program fragment that reads X (by reading each X_i) is executed between the write of 0 to X_1 and the write of 1 to X_2 , then the incorrect value 12 will be obtained for X (i.e., $X_0 = X_1 = 0$ and $X_2 = X_3 = 1$). Thus, this

simple implementation is not serializable. \square

Terminology: In order to avoid confusion, we henceforth refer to the variables of an implemented program as *higher-level* variables, and the variables of an implementation as *lower-level* variables. Note that each higher-level variable corresponds to some set of lower-level variables. We capitalize terms such as “Read” and “Write” when they apply to higher-level variables, and leave them uncapitalized when they apply to lower-level variables. \square

In the above example, the higher-level variable X corresponds to the set of lower-level variables $\{X_0, X_1, X_2, X_3\}$. The assignment statement $X := 4$ “Writes” the value 4 to X . The assignment statement $X_2 := 1$ “writes” the value 1 to X_2 .

Definition: Consider a program fragment that replaces an assignment statement S of the implemented program. We say that the program fragment *Reads* (respectively, *Writes*) each higher-level variable that is read (respectively, written) by the assignment statement S ; the set of all such higher-level variables comprise the *Read Set* (respectively, *Write Set*) of the program fragment. A program fragment is only allowed to access the lower-level variables that correspond to the higher-level variables in its Read Set and Write Set. \square

Consider, for example, the program fragment for $X := 4$ given earlier. The Write Set for the program fragment consists of the single variable X , while the Read Set for the program fragment is empty. Note that the program fragment does not actually write X , but rather writes the corresponding lower-level variables X_0, X_1, X_2 , and X_3 ; variable X exists only on a logical basis in the implementation.

Each execution of a program fragment has the effect of Reading a value from each variable in the Read Set of the program fragment, and Writing a value to each variable in the Write Set of the program fragment. The value “Read from” or “Written to” a given higher-level variable can be defined formally by specifying a function over the values read from and written to the lower-level variables that correspond to that higher-level variable. For instance, in the example considered at the beginning of this section, “the value Written to X ” is defined by the function $v_0 + 2v_1 + 4v_2 + 8v_3$, where each v_i is the value written to X_i . Thus, associated with each program fragment is a set of functions, one for each higher-level variable in the program fragment’s Read Set and Write Set. Note that if a variable appears in both the Read Set and Write Set of a program fragment, then two functions must be specified for that variable, one that defines the value Read from that variable, and another that defines the value Written to that variable.

We now define several concepts that are needed in order to formally define what it means for an implementation to be wait-free, sequentially correct, and serializable. The following definitions apply to any concurrent program.

Definition: A *state* is an assignment of values to the variables of the program. One or more states are designated as *initial states*. \square

Definition: An *event* is an execution of a statement of a program. (In the case of a **do** or **if** statement, “execution of a statement” means the evaluation of each guard in the statement’s set of guards, and the subsequent transfer of control.) \square

Definition: Let t and u be any two states. If u is reached from t via the occurrence of event e , then we say that e is *enabled* at state t and we write $t \xrightarrow{e} u$. \square

Definition: A *history* is a sequence $t_0 \xrightarrow{e_0} t_1 \xrightarrow{e_1} \dots$. Unless noted otherwise, the first state t_0 is assumed to be an initial state. \square

A given statement may be executed many times in a history. Each such statement execution corresponds to a distinct event.

Definition: Event e *precedes* another event f in a history iff e occurs before f in the history. \square

The remaining definitions apply to an implementation.

Definition: Let $S_0; \text{do } B_1 \rightarrow S_1 \parallel \dots \parallel B_N \rightarrow S_N \text{ od}$ be a program fragment of the implementation that replaces some assignment statement of the implemented program. Let h be a history of the implementation. The set of events in h corresponding to an execution of the program fragment — starting with an event that occurs as a result of executing statement S_0 , and followed by all events that subsequently occur as a result of executing the statements in $\text{do } B_1 \rightarrow S_1 \parallel \dots \parallel B_N \rightarrow S_N \text{ od}$ — is called an *operation*. \square

Definition: An operation p *precedes* another operation q in a history iff each event of p precedes all events of q . \square

Observe that the precedes relation is an irreflexive total order on events and an irreflexive partial order on operations.

We now formally define the notions of wait-freedom, sequential correctness, and serializability. These concepts are defined with respect to the histories of an implementation.

Definition: A history of an implementation is *wait-free* iff each operation in the history consists of a finite number of events. \square

Next, we define what it means for a history of an implementation to be “well-formed.” In a well-formed history, the values Read from and Written to the higher-level variables by each operation are well-defined. In defining sequential correctness and serializability, it is sufficient to consider only well-formed histories.

Definition: A history of an implementation is *well-formed* iff each operation in the history is complete, i.e., the execution of the corresponding program fragment terminates. \square

Definition: Let p be an operation in a well-formed history of an implementation, and let $x_1, \dots, x_K := E_1, \dots, E_K$ denote the assignment statement of the implemented program that corresponds to p . Let v_i denote the value Written to x_i by p , and let V_j denote the value obtained by replacing each variable appearing in E_j by the value Read from that variable by p . We say that p is *sequentially correct* iff the condition $v_i = V_i$ holds for each i . A well-formed history of an implementation is *sequentially correct* iff each operation in the history is sequentially correct. \square

Next, we define what it means for a history of an implementation to be serializable. According to this definition, if several operations are executed concurrently, then the net effect should be equivalent to some serial order. Our definition of serializability is similar to the definition of linearizability given by Herlihy and Wing in [13].

Definition: Let h be a well-formed history of an implementation. History h is *serializable* iff the precedence relation on operations (which is a partial order) can be extended³ to a total order \sqsubset where for each operation p in h and each variable x that is Read by p , the value Read by p from x is the same as the last-Written value according to \sqsubset . More specifically,

- if there exists an operation q that Writes x such that $q \sqsubset p \wedge \neg(\exists q' : q' \text{ Writes } x : q \sqsubset q' \sqsubset p)$, then the value Read by p from x is the same as the value Written by q to x ;
- if no such q exists, then the value Read by p from x is the same as the initial value of x as defined in the implemented program. \square

We are now in a position to formally state the correctness condition for an implementation.

Definition: An implementation is *correct* iff each of its histories is wait-free, and each of its well-formed histories is sequentially correct and serializable. \square

5 Necessity Results

In this section, we establish the necessity of Exact-Writing and Single-Phase. In particular, we show that for each of these constraints there exists a nonatomic program violating that constraint, but satisfying the other two constraints. This implies that if we remove either of these two constraints from our set of constraints, then we can no longer make the claim that any program satisfying all of our atomicity constraints is atomic.

The above-mentioned impossibility results are proved by considering the problem of two-process leader election; we call our version of this problem the *binary election problem*. We first show that this problem cannot be solved by a primitive program. We then give two programs that solve this problem: the first program violates Exact-Writing, and the second program violates Single-Phase. It follows, then, that neither of these programs can be implemented by a primitive program.

We now define the binary election problem.

Binary Election Problem: We are required to construct a program consisting of two processes.

³A relation R over a set S *extends* another relation R' over S iff for each a and b in S , $aR'b \Rightarrow aRb$.

```

shared var X, Y, Z : 0..2

initialization X = Y = Z = 0

process p
private var
    electp : boolean; /* decision variable */
    c, d : 0..2
begin
    0: X, Y := 1, 1;
    1: c, d := Y, Z;
    2: if (c, d) = (1, 2) → electp := true
       [] (c, d) ≠ (1, 2) → electp := false
       fi
end

process q
private var
    electq : boolean; /* decision variable */
    e, f : 0..2
begin
    3: Y, Z := 2, 2;
    4: e, f := X, Y;
    5: if (e, f) = (1, 2) → electq := true
       [] (e, f) ≠ (1, 2) → electq := false
       fi
end

```

Figure 2: Proof of Theorem 4.

Each process has a private boolean *decision variable* that it writes only once. The program is required to satisfy the following three conditions.

- *Wait-Freedom*: All histories of the program are finite.
- *Equity*: For each process, there exists a history with a final state in which the value of that process's decision variable is true.
- *Validity*: In the final state of every history, the value of one of the decision variables is true and the value of the other decision variable is false. □

In [5], we prove that the binary election problem cannot be solved by any program satisfying constraints (4) and (5) in the definition of primitive programs. In particular, we show that if such a program satisfies both Equity and Validity, then it has an infinite history, violating Wait-Freedom. Similar proofs have been obtained independently by Chor, Israeli, and Li [9] and by Herlihy [12]. This establishes the following theorem.

Theorem 3: The binary election problem cannot be solved by a primitive program. □

We now show that the binary election problem can be solved if either Exact-Writing or Single-Phase is violated.

Theorem 4: The binary election problem can be solved by a program that violates Exact-Writing.

Proof: Consider the program given in Figure 2. Statement 0 of process p and statement 3 of process q together constitute a violation of Exact-Writing. In particular, it is impossible to partition the variables X , Y , and Z into classes such that each of these two statements writes every variable from one class and no variable from any other class. This program satisfies Wait-Freedom because




Figure 3: Possible values of (X, Y, Z) .

```

shared var X : boolean

process p
private var
  electp : boolean /* decision variable */
begin
  X, electp := ¬X, X
end

process q
private var
  electq : boolean /* decision variable */
begin
  X, electq := ¬X, X
end
  
```

Figure 4: Proof of Theorem 5.

it is loop-free. To see that the program satisfies Equity and Validity, consider Figure 3. This figure depicts the possible values of (X, Y, Z) ; each arrow is labeled by the statement that causes the state change. Based on this figure, we conclude that either the final value of the pair (c, d) equals $(1, 2)$ and the final value of the pair (e, f) differs from $(1, 2)$, or vice versa. This implies that Equity and Validity are satisfied. \square

Theorem 5: The binary election problem can be solved by a program that violates Single-Phase.

Proof: Consider the program given in Figure 4. Note that the assignment statement of each process both reads and writes the shared variable X and hence violates Single-Phase. It is easy to see that this program solves the binary election problem. \square

6 Sufficiency Results

In this section, we discuss the proof strategy for establishing the sufficiency of our atomicity constraints. We prove that any program satisfying the three atomicity constraints is atomic; that is, such a program can be implemented by a primitive program (as described in Section 4). The implementation is facilitated by considering a shared data object, called a *composite register* [2, 3]. A composite register is an array-like variable that consists of a number of components. An operation




Figure 5: $K/L/M/N$ composite register structure.

of a composite register either writes a value to one of the components, or reads the values of all of the components. For convenience, we designate the structure of a given composite register by a four-tuple $K/L/M/N$, where K is the number of components, L is the number of bits per component, M is the number of writers per component, and N is the number of readers. The structure of a $K/L/M/N$ composite register is shown in Figure 5. Note that this figure only depicts the writers for component i .

We now present two theorems that reduce our proof obligation in establishing the sufficiency of our atomicity constraints to one involving only composite registers. First, we show that if the processes of a program communicate only by means of $1/1/1/1$ composite registers, then the program is primitive. Second, we show that any program satisfying our atomicity constraints can be implemented by a program in which processes communicate only by means of $K/L/M/N$ composite registers. It follows from these two theorems that in order to establish the sufficiency of our atomicity constraints, it suffices to prove that a $K/L/M/N$ composite register can be constructed from $1/1/1/1$ composite registers.

Theorem 6: If each shared variable of a program is a $1/1/1/1$ composite register, then the program is primitive.

Proof: A $1/1/1/1$ composite register is a single-reader, single-writer, shared bit. Recall that a program is primitive if its processes communicate only by means of such bits. \square

Theorem 7: Any program satisfying our proposed atomicity constraints can be implemented by a program in which each shared variable is a component of a $K/L/M/N$ composite register.

Proof Sketch: Consider a program P that satisfies the three atomicity constraints: Exact-Writing, Single-Phase, and Finite-Reading. Program P can be implemented by a program Q in which no assignment statement both reads and writes shared variables. In particular, suppose that the following assignment statement appears in some process p of P .

$$x_1, \dots, x_N := E_1, \dots, E_N$$

If this statement both reads a shared variable and writes a shared variable, then, by Single-Phase, it

reads no shared variable that is written by a process other than p . This implies that we can replace this statement by the two statements

$$\begin{aligned} v_1, \dots, v_N &:= E_1, \dots, E_N; \\ x_1, \dots, x_N &:= v_1, \dots, v_N \end{aligned}$$

where each v_i is a new private variable of process p . By similarly replacing each assignment statement of P that both reads and writes shared variables, we obtain program Q .

Now, we show that Q can be implemented by a program R in which each assignment statement writes at most one shared variable. Consider an assignment statement of Q that writes several shared variables Z_1, \dots, Z_K . By Exact-Writing, each Z_i is an unsubscripted variable. Furthermore, if an assignment statement writes some Z_i , then it writes each Z_i and no other shared variable. We can therefore combine Z_1, \dots, Z_K into a new “compound” variable $ZNEW$ as follows:

$$ZNEW : \text{record } Z_1 : \text{type}_1; \dots Z_K : \text{type}_K \text{ end}$$

where type_i is the type of Z_i in Q . (Note that $ZNEW$ is really just an integer or subrange variable. We are using the record type only as “syntactic sugar” to indicate that the bits of $ZNEW$ are partitioned into various fields.) Each assignment statement of Q that writes the K variables Z_1, \dots, Z_K is replaced by an assignment statement that writes the single variable $ZNEW$. We also have to replace each assignment statement that reads some Z_i . For example, the statement

$$a, b := Z_1 + Z_2, Z_3 + Z_4$$

(where a and b are private variables) would be replaced by the statement

$$a, b := ZNEW.Z_1 + ZNEW.Z_2, ZNEW.Z_3 + ZNEW.Z_4$$

Program R is obtained by making similar replacements for each class of shared variables given by Exact-Writing.

From the above construction, if an assignment statement of program R accesses a shared variable, then it either writes a single shared variable or reads one or more shared variables. Moreover, by the Finite-Reading constraint, if an assignment statement of R accesses an element of an infinite shared array, then it either reads that element or writes that element, and accesses no other shared variable.

Based upon these observations, we now show that R can be implemented by a program in which each shared variable is a component of a composite register. The implementation is obtained as follows: each element of an infinite shared array in R is replaced by a composite register with only one component; the remaining shared variables of R (other than elements of infinite shared arrays) are replaced by a single composite register, with each variable corresponding to one component of the register. From the construction of program R , it follows that each assignment statement can be implemented by either a write operation of one component of a composite register, or a read operation of all components of a composite register. \square

Example: We show that the example program considered in Section 3 can be implemented by a program with processes that communicate only via a single composite register. The corresponding

```

type Ctype = record  VW : integer;  X : 1..N;  YZ : boolean  end
shared var C : array[1..2] of Ctype /* composite register */

initialization (C[2].VW = default value)  $\wedge$  (C[1].YZ = C[2].YZ)

process controller
private var request, reply : integer;
           dev : 1..N;
           pending, y : boolean;
           c : array[1..2] of Ctype
begin
  do true  $\rightarrow$  read c := C;
    reply, pending := c[2].VW, (c[1].YZ  $\neq$  c[2].YZ);
    if  $\neg$ pending  $\rightarrow$  consume reply;  produce next request, dev;
      read c := C;  y :=  $\neg$ c[1].YZ;
      write C[1] := (request, dev, y)
    [] pending  $\rightarrow$  skip
  fi
od
end

process device(j : 1..N)
private var request, reply : integer;
           dev : 1..N;
           pending, z : boolean;
           c : array[1..2] of Ctype
begin
  do true  $\rightarrow$  read c := C;
    request, dev, pending := c[1].VW, c[1].X, (c[1].YZ  $\neq$  c[2].YZ);
    if pending  $\wedge$  dev = j  $\rightarrow$  consume request;  produce reply;
      read c := C;  z :=  $\neg$ c[2].YZ;
      write C[2] := (reply, 1, z)
    []  $\neg$ pending  $\vee$  dev  $\neq$  j  $\rightarrow$  skip
  fi
od
end

```

Figure 6: Example program, revisited.

program is shown in Figure 6. There is a single composite register C that replaces the five shared variables V , W , X , Y , and Z of the original program. C has two components, the first of which is written by the *controller*, and the second of which is written by the *devices*.

To explain the structure of C , we return to the original program given in Figure 1. As suggested by Exact-Writing, we partition the five shared variables of this program into two classes $\{V, X, Y\}$ and $\{W, Z\}$. Each of these classes corresponds to one component of C : the class $\{V, X, Y\}$ corresponds to the first component of C , namely $C[1]$, and the class $\{W, Z\}$ corresponds to the second component of C , namely $C[2]$. Note that the class $\{V, X, Y\}$ contains an integer variable, a variable of type $1..N$, and a boolean variable, while the class $\{W, Z\}$ contains an integer variable and a boolean variable. To accommodate both classes, we define each component of C to include a field VW of type integer, a field X of type $1..N$, and a field YZ of type boolean. The correspondence between the shared variables of the original program and the composite register C is then as follows: variable V corresponds to $C[1].VW$, variable W corresponds to $C[2].VW$, variable X corresponds to $C[1].X$, variable Y corresponds to $C[1].YZ$, and variable Z corresponds to $C[2].YZ$. Note that $C[2].X$ is unused.

Using our tuple notation, C would be classified as a $2/\infty/N/N + 1$ composite register: it has two components; each component consists of an infinite number of bits (since each VW field is an integer); each component can be written by N processes (although the *controller* actually only writes the first component and the N *devices* the second); and all $N + 1$ processes can read C . Note that in Figure 6 we have annotated the program with the keywords `read` and `write` to emphasize those assignment statements that access the composite register C . This convention will also be adopted in the remainder of the paper. \square

It follows from these two theorems that in order to establish the sufficiency of our atomicity constraints, our proof obligation is reduced to the following.

Proof Obligation: Show that a $K/L/M/N$ composite register can be implemented in terms of $1/1/1/1$ composite registers. \square

Fortunately, this proof obligation can be divided into a series of steps, one for each of the parameters K , L , M , and N . Constructions for these steps can be found in the Ph.D. dissertation of the first author [4], as well as in the literature. The steps involved in constructing a $K/L/M/N$ composite register from $1/1/1/1$ composite registers are illustrated in Figure 7. Each arrow in this figure is labeled by a reference to the paper(s) in which the corresponding register construction(s) appear.

One of the steps shown in Figure 7 is considered in detail in Section 8. In particular, we present a construction of a $1/L/1/1$ composite register from $1/1/1/1$ composite registers. This construction is included to give the reader (of this paper) some idea as to the structure of a composite register construction. Our $1/L/1/1$ construction is the first that we know of that explicitly addresses the case in which L is allowed to be infinite.

In proving that a composite register construction is correct, the principal difficulty that arises is that of establishing that all well-formed histories of the construction are serializable. The definition




Figure 7: Constructing a $K/L/M/N$ composite register from $1/1/1/1$ composite registers.

of serializability given in Section 4, while intuitive, is rather difficult to use directly. In the next section, we present a lemma that gives a set of conditions that are sufficient for establishing that a well-formed history of a construction is serializable. Intuitively, a history is serializable if each operation in the history can be shrunk to a point; that is, there exists a point between the first and last events of each operation at which the operation appears to take effect. For this reason, this lemma is referred to as the “Shrinking Lemma.”

7 Shrinking Lemma

Before presenting the Shrinking Lemma, we first describe the basic structure of a composite register construction. We define a construction in a rather abstract way: we focus only on the Read and Write operations of the constructed register, while ignoring the processes that invoke these operations. More specifically, a construction consists of a set of Writer procedures and a set of Reader procedures that communicate via a set of lower-level variables. A Writer procedure is invoked by a process in order to Write a value to a component of the constructed composite register. A Reader procedure is invoked by a process in order to Read the values of all of the components of the constructed composite register. Each Writer procedure has one input parameter indicating the value to be

Written; each Reader procedure has one output parameter for each component of the constructed register.

The Writer and Reader procedures of a construction correspond to the program fragments introduced in Section 4. The Read Set of a Writer procedure consists of the single input parameter of that procedure, while its Write Set consists of the single component Written by that procedure. The Read Set of a Reader procedure includes all of the components of the constructed register, while its Write Set includes all of its output parameters. In a history of a construction, the Writer and Reader procedures are invoked in an arbitrary manner.

The correctness condition for a construction is the same as defined in Section 4. That is, a construction is correct iff each of its histories is wait-free, and each of its well-formed histories is sequentially correct and serializable. Establishing that a history is wait-free and sequentially correct is usually straightforward. The major difficulty lies in establishing that a history is serializable, so it is this requirement that we will concentrate on here.

It is rather difficult to prove that a history of a composite register construction is serializable by appealing directly to the definition of serializability given in Section 4. The Shrinking Lemma, given below, reduces our proof obligation to five distinct conditions. Before stating this lemma, we first introduce some terminology.

Definition: A Write operation of component k of the constructed composite register, where $0 \leq k < K$, is called a k -Write operation. \square

In order to avoid special cases when proving the correctness of a construction, we make the following assumption concerning the initial Write operations.

Initial Writes: For each k , where $0 \leq k < K$, there exists a k -Write operation that precedes each other k -Write operation and all Read operations. \square

Shrinking Lemma: A well-formed history h of a $K/L/M/N$ composite register construction is serializable if for each k , where $0 \leq k < K$, there exists a function ϕ_k that maps every Read operation and k -Write operation in h to some natural number, such that the following five conditions hold.

- *Uniqueness:* For each pair of distinct k -Write operations v and w in h , $\phi_k(v) \neq \phi_k(w)$. Furthermore, if v precedes w , then $\phi_k(v) < \phi_k(w)$.
- *Integrity:* For each Read operation r in h , and for each k in the range $0 \leq k < K$, there exists a k -Write operation w in h such that $\phi_k(r) = \phi_k(w)$. Furthermore, the value Read by r for component k is the same as the value Written by w .
- *Proximity:* For each Read operation r in h and each k -Write operation w in h , if r precedes w then $\phi_k(r) < \phi_k(w)$, and if w precedes r then $\phi_k(w) \leq \phi_k(r)$.
- *Read Precedence:* For each pair of Read operations r and s in h , if $(\exists k :: \phi_k(r) < \phi_k(s))$ or if r precedes s , then $(\forall k :: \phi_k(r) \leq \phi_k(s))$.
- *Write Precedence:* For each Read operation r in h , and each j -Write operation v and k -Write

operation w in h , where $0 \leq j < K$ and $0 \leq k < K$, if v precedes w and $\phi_k(w) \leq \phi_k(r)$, then $\phi_j(v) \leq \phi_j(r)$. \square

Note that for $K = 1$, Write Precedence is a direct consequence of Uniqueness.

The correctness proof for the Shrinking Lemma is given in Appendix 1. The proof is somewhat tedious, but is not hard. The proof strategy is as follows. We first augment the precedence relation on operations in history h by adding pairs of operations. These added pairs of operations are defined based on the five conditions of the lemma. We then show that the resulting relation is an irreflexive partial order, i.e., it is irreflexive and transitive. Finally, we show that any extension of this relation to an irreflexive total order satisfies the conditions in the definition of serializability given in Section 4.

8 $1/L/1/1$ Construction

In this section, we establish that a $1/L/1/1$ composite register can be constructed from $1/1/1/1$ composite registers (even if L is infinite). The construction is described informally below, and its correctness proof is given in Appendix 2.

The construction is depicted in Figure 8. (We assume that any arithmetic expression involving L in this figure is replaced by ∞ for the case in which L is infinite.) We begin by giving a brief description of how the construction works; a more detailed description is given below. The Reader and the Writer communicate by means of four “buffers.” Each buffer is an array of bits. The Writer Writes a value to the constructed register by copying the binary representation of this value into one of the four buffers. The Reader Reads a value from the constructed register by scanning one of the four buffers. To ensure that the Reader obtains a correct value, the Reader and the Writer are coordinated so that they never access the same buffer at the same time. This technique has been employed in several other similar constructions; see, for example, Burns and Peterson [8], Kirousis et al. [15], and Tromp [27].

The following shared variables are used in the construction.

$Y[i, j]$: A buffer that is used to store the binary representation of a value as Written by a Write operation. There are four such buffers since i and j each range over $\{0, 1\}$. Each buffer is an array of bits that is written by the Writer and read by the Reader. If L is finite, then $Y[i, j]$ has $2L + 1$ elements, and if L is infinite, then $Y[i, j]$ has an infinite number of elements. (Note that Y is really a three-dimensional array. However, in describing how the construction works, we find it convenient to view Y as a two-dimensional array, where each element is itself an array.)

$Q[i]$: An integer auxiliary variable that is written by the Writer and read by the Reader. There are two such variables since i ranges over $\{0, 1\}$. Note also that the Reader and Writer each have a private auxiliary variable called *id*. These auxiliary variables have been introduced solely to facilitate the proof of correctness.

The following shared variables are used to ensure that the Reader and Writer never access the same buffer at the same time. This is described in detail below.

```

type Ytype = array[0..2L] of 0..1
shared var Y : array[0..1][0..1] of Ytype;
        Z : array[0..1] of 0..1;
        Q : array[0..1] of integer; /* auxiliary variable */
        WP, RP : 0..1

initialization WP = RP  $\wedge$  ( $\forall i, j : 0 \leq i \leq 1 \wedge 0 \leq j \leq 1 : (\exists k : k \geq 0 : Y[i, j][2k] = 0)$ )

procedure Reader returns valtype
private var
    id : integer; /* auxiliary variable */
    val : valtype;
    ralt, rp, x, y : 0..1;
    k : 0..L + 1
initialization
    rp = RP
begin
    0: read rp := WP;
    1: write RP := rp;
    2:  $\langle$  read ralt := Z[rp]; read id := Q[rp]  $\rangle$ ;
    3: val, k := 0, 0;
    4: do k  $\leq$  L  $\rightarrow$ 
        read x := Y[rp, ralt][2k];
        if x = 0  $\rightarrow$  exit [] x  $\neq$  0  $\rightarrow$  skip fi;
        read y := Y[rp, ralt][2k + 1];
        val, k := val + y  $\cdot$  2k, k + 1
    od;
    5: return(val)
end

procedure Writer(val : valtype)
private var
    id : integer; /* auxiliary variable */
    d, wp, walt : 0..1;
    n : 1..L;
    k : 0..L
initialization
    id = 0  $\wedge$  d = wp = WP
begin
    0: id := id + 1;
    1: read d := RP;
    2: wp, walt, n := d  $\oplus$  1, walt  $\oplus$  1, number of bits in val;
    3: write Y[wp, walt][2n] := 0;
    4: k := 0;
    5: do k < n  $\rightarrow$ 
        write Y[wp, walt][2k] := 1;
        write Y[wp, walt][2k + 1] := val[k];
        k := k + 1
    od;
    6:  $\langle$  write Z[wp] := walt; write Q[wp] := id  $\rangle$ ;
    7: write WP := wp
end

```

Figure 8: 1/L/1/1 construction.

$Z[i]$: A bit that is written by the Writer and read by the Reader. There are two such bits since i ranges over $\{0, 1\}$.

WP : A modulo-2 “write pointer” that is written by the Writer and read by the Reader. (We use \oplus to denote modulo-2 addition.)

RP : A modulo-2 “read pointer” that is written by the Reader and read by the Writer.

To simplify the Reader and Writer procedures in Figure 8, we assume that the private variables of each procedure retain their values between invocations. Observe that $Z[i]$ and the auxiliary variable $Q[i]$ are written (read) by the Writer (Reader) in a single action. This is denoted by the angle brackets ‘ \langle ’ and ‘ \rangle ’.

A Write operation Writes a new value to the constructed register by copying the binary representation of this value into the buffer $Y[wp, walt]$. Thus, wp and $walt$ together constitute a two-bit “pointer” to one of the four shared buffers. The value of wp is obtained by first reading RP and then incrementing the value read. The value of $walt$ is obtained by incrementing the value assigned to $walt$ by the previous Write operation. After updating $Y[wp, walt]$, the value of $walt$ is made available to the Reader by writing it to $Z[wp]$, and the value of wp is made available to the Reader by writing it to WP .

Each pair of bits ($Y[wp, walt][2k], Y[wp, walt][2k + 1]$), where $k \geq 0$, encodes a single bit of the Written value. The pair $(1, 0)$ encodes a bit with value 0 and the pair $(1, 1)$ encodes a bit with value 1. The pair $(0, -)$ is used to indicate the end of a string of bits. Note that such a string can be arbitrarily long since L may be infinite — hence the need for the encoding.

A Read operation Reads a value from the constructed register by scanning one of the four shared buffers. A Read operation scans the buffer that is pointed to by the pair of bits rp and $ralt$. The value of rp is obtained by reading WP and the value of $ralt$ is obtained by reading $Z[rp]$. The value of rp is made available to the Writer by writing it to RP .

The correctness of this construction follows from the following fact (which is proved as a lemma in Appendix 2): if a Read operation executes its `do` statement while a Write operation executes its `do` statement, then $rp \neq wp$ or $ralt \neq walt$. This implies that no Write operation interferes with a Read operation as it reads from one of the four shared buffers.

We now compute the complexity of our $1/L/1/1$ construction. The space complexity of a construction is determined by the number of shared $1/1/1/1$ composite registers required. If L is infinite, then an infinite number of $1/1/1/1$ registers are obviously required. So, assume that L is finite. Then, the complexity of each of the shared variables (other than the auxiliary variable $Q[i]$) is as follows.

- $Y[i, j]$, where $0 \leq i, j \leq 1$, uses $2L + 1$ bits.
- $Z[i]$, where $0 \leq i \leq 1$, uses 1 bit.
- WP uses 1 bit.
- RP uses 1 bit.

Therefore, the space complexity is $8L + 8$. The lower bound to construct an L -bit register is $O(L)$. Thus, our construction is asymptotically optimal.

The time complexity of a $1/L/1/1$ construction is determined by the number of reads and writes of shared $1/1/1/1$ registers required to Read and Write the constructed register. The time complexity of a Read in our $1/L/1/1$ construction is $O(L')$, where $L' < L$ is the size (in bits) of the value Read. The time complexity of a Write is $O(L')$, where $L' < L$ is the size (in bits) of the value Written.

9 Concluding Remarks

We have presented a new criterion for the atomicity of concurrent programs along with a simple method for checking whether a given program is atomic. This method consists of three constraints, which can often be checked based on the syntax of a particular program. The class of programs

considered atomic according to our atomicity criterion is larger than the class considered atomic according to the traditional Owicki-Gries atomicity criterion.

Our atomicity criterion is motivated by the following observation: Interleaving semantics is a suitable abstraction only if we restrict our attention to programs with sufficiently simple actions. According to our criterion, an action is “sufficiently simple” if it can be implemented in a wait-free, serializable manner in terms of primitive actions. (In [11], Gouda argues that we should also restrict our attention to only certain kinds of assertions when reasoning about the behavior of a program under interleaving semantics. In particular, he considers a restricted notion of parallelism and argues that for interleaving semantics to be a suitable abstraction for this notion of parallelism, we should limit our attention to “parallelizable” assertions.)

Our notion of serializability is equivalent to the definition of linearizability given by Herlihy and Wing in [13], for the special case of implementing a program in which shared variables are accessed only by assignment statements. A similar notion has also been proposed by Misra [22]. In particular, Misra gives a set of axioms for constructing concurrent hardware registers, and shows that if a register satisfies these axioms, then concurrent accesses to the register can be viewed as being interleaved.

To establish the sufficiency of our atomicity constraints, we considered in Section 6 a shared data object called a composite register. The notion of a composite register is an extension of an atomic register, and was first introduced by Anderson in [2, 3]. A composite register is equivalent to the atomic snapshot primitive defined by Afek et al. in [1]. The notion of an atomic register was first defined by Lamport [16]. An atomic register is a composite register with only one component.

We should point out that it may be possible to relax the Exact-Writing and Single-Phase constraints somewhat, despite the necessity results proved in Section 5. For example, according to Exact-Writing, if two assignment statements in different processes write a common shared variable, then they write the same set of shared variables. In establishing the necessity of Exact-Writing, we considered in Theorem 4 a program in which the two assignment statements $X, Y := 1, 1$ and $Y, Z := 2, 2$ appeared in different processes. Each of these assignment statements writes one shared variable that is written by the other, and one that is not. This necessity proof leaves open the possibility of relaxing Exact-Writing to allow an assignment statement of one process to write a subset of the variables that are written by an assignment of another process; for example, $X, Y, Z := 1, 1, 1$ and $Y, Z := 2, 2$.

It may also be possible to relax the atomicity constraints based upon the semantics of a particular program. For instance, consider a program in which one process executes the statement $X := Y$ and another process executes the statement $Y := X$. This program violates Single-Phase. However, if the two processes are synchronized in such a way so that these two statements are never enabled at the same time, then it may be reasonable to consider both statements to be atomic. We have tried as much as possible to avoid the need for analyzing the semantics of a program when appealing to our atomicity constraints.

As mentioned earlier, the version of Exact-Writing considered in this paper is overly restrictive because it allows an assignment statement to write several shared variables only if all such variables are unsubscripted. In [4], a more permissive version of Exact-Writing is considered that allows an assignment statement to write several subscripted variables. This version of Exact-Writing is similar

to the one given in this paper, except for the partitioning of shared variables into classes. In [4], a class is allowed to consist of several shared variables in the following cases.

- A class may consist of a number of unsubscripted variables.
- A class may consist of all elements of a finite array.
- If A, B, \dots, C are single-dimensional arrays whose subscripts have the same range, then their elements can be partitioned by subscript, i.e., $A[i], B[i], \dots, C[i]$ is a class for each i . Multi-dimensional arrays may be similarly partitioned.

The above version of Exact-Writing is still, in fact, more restrictive than is necessary. However, there is a tradeoff here: the more permissive we are concerning classes, the more difficult it will be to establish the sufficiency of our atomicity constraints.

We leave the question of whether Finite-Reading is necessary as an open problem. To prove that Finite-Reading is necessary, we are required to demonstrate the existence of a nonatomic program that violates Finite-Reading, but satisfies Exact-Writing and Single-Phase. On the other hand, to prove that Finite-Reading is not necessary (i.e., can be removed from our proposed set of atomicity constraints), we are required to show that any program satisfying only Exact-Writing and Single-Phase can be implemented in terms of single-reader, single-writer, shared bits. Note that the approach that we took in Theorem 2 will not work in this case. The principal difficulty that arises is illustrated by the following example. Consider a concurrent program with a process that includes the following program fragment, where the infinite array $X[0..\infty]$ is shared.

```

i := 0;
do true → a[i], b[i], i := X[i], X[i + 1], i + 1 od

```

This program violates Finite-Reading because each assignment statement $a[i], b[i], i := X[i], X[i + 1], i + 1$, where $i \geq 0$, reads two elements of the infinite shared array $X[0..\infty]$. In order to implement the variables of this program in terms of composite registers, it is necessary to partition the elements of X into finite classes and implement each such class by using a single composite register. This partitioning is necessary because a composite register, by definition, consists of a finite number of components. (An atomic snapshot of an infinite number of components cannot be implemented in a wait-free manner.) However, it is impossible to partition the elements of X so that all shared variables accessed by each assignment statement come from the same class.

Acknowledgements: We would like to thank Ted Herman and Fred Schneider for their comments on an earlier draft of this paper. We would also like to thank the following members of the UT Distributed Systems Discussion Group for their careful reading of a draft of this paper: Yeturu Aahlad, Bryan Bayerdorffer, Ken Calvert, Al Carruth, Edgar Knapp, Jacob Kornerup, David Naumann, and J.R. Rao.

Appendix 1: Proof of the Shrinking Lemma

Shrinking Lemma: A well-formed history h of a $K/L/M/N$ composite register construction is serializable if for each k , where $0 \leq k < K$, there exists a function ϕ_k that maps every Read operation and k -Write operation in h to some natural number, such that the following five conditions hold.

- *Uniqueness*: For each pair of distinct k -Write operations v and w in h , $\phi_k(v) \neq \phi_k(w)$. Furthermore, if v precedes w , then $\phi_k(v) < \phi_k(w)$.
- *Integrity*: For each Read operation r in h , and for each k in the range $0 \leq k < K$, there exists a k -Write operation w in h such that $\phi_k(r) = \phi_k(w)$. Furthermore, the value Read by r for component k is the same as the value Written by w .
- *Proximity*: For each Read operation r in h and each k -Write operation w in h , if r precedes w then $\phi_k(r) < \phi_k(w)$, and if w precedes r then $\phi_k(w) \leq \phi_k(r)$.
- *Read Precedence*: For each pair of Read operations r and s in h , if $(\exists k :: \phi_k(r) < \phi_k(s))$ or if r precedes s , then $(\forall k :: \phi_k(r) \leq \phi_k(s))$.
- *Write Precedence*: For each Read operation r in h , and each j -Write operation v and k -Write operation w in h , where $0 \leq j < K$ and $0 \leq k < K$, if v precedes w and $\phi_k(w) \leq \phi_k(r)$, then $\phi_j(v) \leq \phi_j(r)$.

Proof: The proof strategy is as follows. We first augment the precedence relation on operations in history h by adding pairs of operations. We then show that the resulting relation is an irreflexive partial order, i.e., it is irreflexive and transitive. Finally, we show that any extension of this relation to an irreflexive total order satisfies the conditions in the definition of serializability given in Section 4.

In the remainder of the proof, we use r and s to denote Read operations in history h , v and w to denote Write operations in h , and x , y , and z to denote arbitrary operations in h . We also assume that i , j , and k each range over $\{0, \dots, K-1\}$. If x precedes y in h , then we write $x \triangleleft y$. We let $(x \trianglelefteq y) \equiv (x = y \vee x \triangleleft y)$. We now define six relations A , B , C , D , E , and F ; in these definitions, we assume that v and v' denote j -Write operations, and w and w' denote k -Write operations.

- A includes all pairs (x, y) such that $x \triangleleft y$.
- B includes all pairs (w, r) such that $\phi_k(w) \leq \phi_k(r)$, and all pairs (r, w) such that $\phi_k(r) < \phi_k(w)$.
- C includes all pairs (r, s) such that $(\exists k :: \phi_k(r) < \phi_k(s))$.
- D includes all pairs (v, w) such that $(\exists r :: vBr \wedge rBw)$.
- E includes all pairs (v, w) , such that $v \neq w$ and for some v' and w' ,

$$\phi_j(v) \leq \phi_j(v') \wedge v' \trianglelefteq w' \wedge \phi_k(w') \leq \phi_k(w) .$$

- $F \equiv A \cup B \cup C \cup D \cup E$

Relation A is the precedence relation on operations in history h . Relation F is obviously an extension of A . We now show that F is irreflexive and transitive. To prove that F is irreflexive, we are obligated to show that $xFy \Rightarrow x \neq y$. If xAy , then because A is an irreflexive partial order, $x \neq y$. If xBy , then one of x and y is a Read operation and the other is a Write operation, so $x \neq y$. If xCy , then $\phi_k(x) < \phi_k(y)$ for some k , which implies that $x \neq y$. If xDy , then because relation B

(which is used to define D) totally orders each Read with respect to all Writes, we have $x \neq y$. If xEy , then by the definition of E , $x \neq y$. Therefore, we conclude that relation F is irreflexive.

In the proof of transitivity, we use the following three properties.

Property 1: For each pair of Read operations r and s , $rFs \Rightarrow (\forall k :: \phi_k(r) \leq \phi_k(s))$.

Proof of Property 1: Assume that rFs holds. Of the five relations that define F , only A and C can relate two Read operations. Therefore, rAs holds or rCs holds. In the former case (i.e., r precedes s), by Read Precedence, $(\forall k :: \phi_k(r) \leq \phi_k(s))$. In the latter case, by the definition of C , $(\exists k :: \phi_k(r) < \phi_k(s))$; hence, by Read Precedence, $(\forall k :: \phi_k(r) \leq \phi_k(s))$. \square

Property 2: For each Read operation r and k -Write operation v , $rFv \Rightarrow \phi_k(r) < \phi_k(v)$ and $vFr \Rightarrow \phi_k(v) \leq \phi_k(r)$.

Proof of Property 2: We prove that $rFv \Rightarrow \phi_k(r) < \phi_k(v)$; the proof that $vFr \Rightarrow \phi_k(v) \leq \phi_k(r)$ is similar. Assume that rFv holds. Of the five relations that define F , only A and B can relate a Read operation and a Write operation. Therefore, rAv holds or rBv holds. In the former case (i.e., r precedes v), by Proximity, $\phi_k(r) < \phi_k(v)$. In the latter case, by the definition of B , $\phi_k(r) < \phi_k(v)$. \square

Property 3: For each pair of Write operation v and w , $vAw \Rightarrow vEw$.

Proof of Property 3: Let v be a j -Write operation and let w be a k -Write operation such that vAw holds. If $j \neq k$, then $v \neq w$. If $j = k$, then by Uniqueness, $\phi_j(v) < \phi_j(w)$, which implies that $v \neq w$. Therefore, letting $v' = v$ and $w' = w$, we have

$$v \neq w \wedge \phi_j(v) \leq \phi_j(v') \wedge v' \triangleleft w' \wedge \phi_k(w') \leq \phi_k(w) .$$

This implies that vEw . \square

To prove that F is transitive, we are obligated to show that $xFy \wedge yFz \Rightarrow xFz$. We have to consider eight cases since each of x , y , and z can be either a Read operation or a Write operation.

Case 1: x , y , and z are all Read operations. Of the five relations that define F , only A and C can relate two Read operations. If xAy and yAz , then because A is a partial order, xAz . Now, suppose that xCy holds. By the definition of C , we have $\phi_j(x) < \phi_j(y)$ for some j . By Property 1, $(\forall k :: \phi_k(y) \leq \phi_k(z))$. Therefore, by transitivity, $\phi_j(x) < \phi_j(z)$, i.e., xCz . Similar reasoning applies if yCz holds.

Case 2: x and y are Read operations and z is a j -Write operation. By Property 1, $(\forall k :: \phi_k(x) \leq \phi_k(y))$. By Property 2, $\phi_j(y) < \phi_j(z)$. Therefore, by transitivity, $\phi_j(x) < \phi_j(z)$, i.e., xBz .

Case 3: x and z are Read operations and y is a j -Write operation. By Property 2, $\phi_j(x) < \phi_j(y)$ and $\phi_j(y) \leq \phi_j(z)$. Therefore, by transitivity, $\phi_j(x) < \phi_j(z)$, i.e., xCz .

Case 4: x is a Read operation, y is a j -Write operation, and z is a k -Write operation. Of the five relations that define F , only A , D , and E can relate two Write operations. Thus, by Property 3, yDz holds or yEz holds. If yDz holds, then by the definition of D , there exists a Read operation r such that yBr and rBz . Because $xFy \wedge yBr$ holds, by Case 3, we have xFr . Because $xFr \wedge rBz$ holds, by Case 2, we have xFz .

Now, consider the case yEz . By the definition of E , there exists a j -Write operation v and a k -Write operation w such that

$$\phi_j(y) \leq \phi_j(v) \wedge v \triangleleft w \wedge \phi_k(w) \leq \phi_k(z) .$$

By Property 2, $\phi_j(x) < \phi_j(y)$; thus, by transitivity, $\phi_j(x) < \phi_j(v)$. If $v = w$ (which implies that $j = k$), then $\phi_k(x) < \phi_k(w)$. If, on the other hand, $v \triangleleft w$, then by the contrapositive of Write Precedence, $\phi_k(x) < \phi_k(w)$. Therefore, by transitivity, $\phi_k(x) < \phi_k(z)$, i.e., xBz .

Case 5: x is a j -Write operation and both y and z are Read operations. By Property 2, $\phi_j(x) \leq \phi_j(y)$. By Property 1, $(\forall k :: \phi_k(y) \leq \phi_k(z))$. Therefore, by transitivity, $\phi_j(x) \leq \phi_j(z)$, i.e., xBz .

Case 6: x is a j -Write operation, y is a Read operation, and z is a k -Write operation. By Property 2, $\phi_j(x) \leq \phi_j(y)$ and $\phi_k(y) < \phi_k(z)$. Hence, by the definition of B , xBz and yBz . Therefore, by the definition of D , $x Dz$.

Case 7: x is a j -Write operation, y is a k -Write operation, and z is a Read operation. Of the five relations that define F , only A , D , and E can relate two Write operations. Thus, by Property 3, $x Dy$ holds or $x Ey$ holds. If $x Dy$ holds, then by the definition of D , there exists a Read operation r such that $x Br$ and $r By$. Because $r By \wedge y Fz$ holds, by Case 3, we have $r Fz$. Because $x Br \wedge r Fz$ holds, by Case 5, we have $x Fz$.

Now, consider the case $x Ey$. By the definition of E , there exists a j -Write operation v and a k -Write operation w such that

$$\phi_j(x) \leq \phi_j(v) \wedge v \triangleleft w \wedge \phi_k(w) \leq \phi_k(y) .$$

By Property 2, $\phi_k(y) \leq \phi_k(z)$. Thus, by transitivity, $\phi_k(w) \leq \phi_k(z)$. If $v = w$ (which implies that $j = k$), then $\phi_j(v) \leq \phi_j(z)$. If, on the other hand, $v \triangleleft w$, then by Write Precedence, $\phi_j(v) \leq \phi_j(z)$. Hence, by transitivity, $\phi_j(x) \leq \phi_j(z)$, i.e., xBz .

Case 8: x , y , and z are all Write operations. Of the five relations that define F , only A , D , and E can relate two Write operations. Thus, by Property 3, $x Dy$ holds or $x Ey$ holds, and $y Dz$ holds or $y Ez$ holds. If $x Dy$ holds, then there exists a Read operation r such that $x Br$ and $r By$. Because $r By \wedge y Fz$ holds, by Case 4, we have $r Fz$. Because $x Br \wedge r Fz$ holds, by Case 6, we have $x Fz$. The case in which $y Dz$ holds is similar.

The remaining possibility is $x Ey$ and $y Ez$. Assume that x is an i -Write operation, y is a j -Write operation, and z is a k -Write operation. By the definition of E , there exist an i -Write operation v , j -Write operations w and v' , and a k -Write operation w' such that

$$x \neq y \wedge \phi_i(x) \leq \phi_i(v) \wedge v \triangleleft w \wedge \phi_j(w) \leq \phi_j(y)$$

and

$$y \neq z \wedge \phi_j(y) \leq \phi_j(v') \wedge v' \trianglelefteq w' \wedge \phi_k(w') \leq \phi_k(z) .$$

There are three possibilities to consider: $i = j$, $j = k$, and $i \neq j \wedge j \neq k$. First, suppose that $i = j$. We show that xEz holds by first proving that $\phi_j(x) < \phi_j(v')$. Because $v \trianglelefteq w$, by Uniqueness, $\phi_j(v) \leq \phi_j(w)$. Therefore, by transitivity, $\phi_j(x) \leq \phi_j(y)$. Because $x \neq y$, Uniqueness implies that $\phi_j(x) \neq \phi_j(y)$. Thus, $\phi_j(x) < \phi_j(y)$. Thus, by transitivity, $\phi_j(x) < \phi_j(v')$.

We now show that xEz . If $i \neq k$, then because x is an i -Write operation and z is a k -Write operation, $x \neq z$. If $i = k$, then by Uniqueness, $\phi_j(v') \leq \phi_j(w')$; thus, by transitivity, $\phi_j(x) < \phi_j(z)$, which implies that $x \neq z$. Therefore, we conclude for the case $i = j$ that

$$x \neq z \wedge \phi_j(x) < \phi_j(v') \wedge v' \trianglelefteq w' \wedge \phi_k(w') \leq \phi_k(z) .$$

Thus, xEz .

The case in which $j = k$ is similar to the case $i = j$.

Now suppose that $i \neq j$ and $j \neq k$. In this case, we prove that xEz by first showing that $v \triangleleft w'$. Because $i \neq j$ and because v is an i -Write operation and w a j -Write operation, we have $v \neq w$. Similarly, because $j \neq k$, we have $v' \neq w'$. Hence, $v \triangleleft w$ and $v' \triangleleft w'$. Observe that, by transitivity, $\phi_j(w) \leq \phi_j(v')$. Therefore, by Uniqueness, $\neg(v' \triangleleft w)$. Also, observe that $v \triangleleft w \wedge v' \triangleleft w' \Rightarrow v' \triangleleft w \vee v \triangleleft w'$. Thus, $v \triangleleft w'$. Hence, the following expression holds.

$$\phi_i(x) \leq \phi_i(v) \wedge v \triangleleft w' \wedge \phi_k(w') \leq \phi_k(z)$$

If $i \neq k$, then because x is an i -Write operation and z a k -Write operation, $x \neq z$. If, on the other hand, $i = k$, then by Uniqueness, $\phi_i(v) < \phi_i(w')$; hence, $\phi_i(x) < \phi_i(z)$, which implies that $x \neq z$. Therefore, we conclude that xEz holds.

Thus, we have established that F is an irreflexive partial order. We now show that any extension of F to an irreflexive total order satisfies the conditions given in the definition of serializability. The following property is used in the proof.

Property 4: For each pair of k -Write operations v and w , $vFw \Rightarrow \phi_k(v) < \phi_k(w)$.

Proof of Property 4: Assume that vFw holds. Then, by Property 3, vDw holds or vEw holds. If vDw holds, then there exists a Read operation r such that $vBr \wedge rBw$. By the definition of B , $\phi_k(v) \leq \phi_k(r)$ and $\phi_k(r) < \phi_k(w)$. This implies that $\phi_k(v) < \phi_k(w)$. If, on the other hand, vEw holds, then there exists k -Write operations v' and w' such that

$$\phi_k(v) \leq \phi_k(v') \wedge v' \trianglelefteq w' \wedge \phi_k(w') \leq \phi_k(w) .$$

By Uniqueness, $\phi_k(v') \leq \phi_k(w')$. This implies that $\phi_k(v) \leq \phi_k(w)$. By the definition of E , $v \neq w$. Therefore, by Uniqueness, $\phi_k(v) < \phi_k(w)$. \square

Let r be a Read operation. By Integrity, there exists a k -Write operation v such that $\phi_k(v) = \phi_k(r)$ and the value Written by v is the same as the value Read by r for component k . By the

definition of B , vFr . Moreover, by Properties 2 and 4, $\neg(\exists w : w \text{ is a } k\text{-Write} : vFw \wedge wFr)$. Observe that, by the definition of B , F totally orders each Read with respect to all Writes. Also, by the definition of E and Uniqueness, the Writes on a given component are totally ordered. Thus, any extension of relation F to an irreflexive total order satisfies the conditions given in the definition of serializability. \square

Appendix 2: Correctness Proof for 1/L/1/1 Construction

In this appendix, we prove that the 1/L/1/1 construction presented in Section 8 is correct by considering an arbitrary history of the construction. We first prove that the given history is wait-free. We then show that if the given history is well-formed, then it is serializable. (Sequential correctness is trivial.) In the proof of serializability, we define a function ϕ for the history, and then prove that the defined ϕ satisfies the Uniqueness, Integrity, Proximity, and Read Precedence conditions of the Shrinking Lemma. (Note that, because $K = 1$, we drop the subscript of ϕ . Also, recall that Write Precedence follows from Uniqueness when $K = 1$.)

The following notational conventions and definitions are used in the rest of this appendix.

Notation: We use r and s to denote Read operations, v and w to denote Write operations, p and q to denote arbitrary operations, and t and u to denote states. We assume that i and j each range over $\{0, 1\}$ and that $k \geq 0$. \square

Notation: In order to avoid using too many parentheses, we define a binding order for the symbols that we use. The following is a list of these symbols, grouped by binding power; the groups are ordered from highest binding power to lowest.

[], ()
 .
 !
 $\neg, :$
 $+, \oplus$
 $=, \neq, <, >, \leq, \geq, \prec, \preceq$
 \wedge, \vee
unless
 \Rightarrow, \equiv
 \models

\square

Definition: If event e precedes event f , then we write $e \prec f$. We let $(e \preceq f) \equiv (e = f \vee e \prec f)$. \square

Definition: Let p be an operation, and let x be any private variable of p . Then, $p!x$ denotes the final value of variable x as assigned by operation p . \square

Definition: Let p be an operation of some Reader or Writer procedure and let i be a label of a

statement in that procedure. We denote the event corresponding to the execution of statement i in operation p by $p:i$. \square

Definition: If E is an expression that holds at state t , then we write $t \models E$. \square

Whenever we say that an expression is true without referring to a particular state, we mean that the expression is an *invariant*, i.e., is true at each state of every history.

Assumption: We assume that each state in every history is distinct. This assumption is easy to ensure by introducing an integer auxiliary variable that is incremented with each event. \square

Definition: Consider the history $t_0 \xrightarrow{e_0} \dots t_i \xrightarrow{e_i} t_{i+1} \dots$. We say that t_i is the state *prior to* the event e_i , t_{i+1} is the state *following* e_i , and e_i is the event *prior to* state t_{i+1} . \square

Note that the event prior to a given state is uniquely defined since, by assumption, each state appears at most once in a given history.

Definition: Let e be the event corresponding to the execution of the statement `read $x := Y$` in an operation p , where x is a private variable and Y is a shared variable. If f is the last event to write Y before e , then we say that f *determines* $p!x$. \square

Definition: Let X be a shared variable of the construction, and let p be an operation. The assertion $last(X) = p$ holds at a state iff the last event to write to X before that state is an event of p . \square

Definition: Let e be an event in some history. Then, $after(e)$ is true at a state iff the state occurs after the event e . \square

Definition: Let A and B be two expressions over the variables of the construction. The assertion A *unless* B holds iff for every pair of consecutive states in any history, if $A \wedge \neg B$ holds in the first state, then $A \vee B$ holds in the second state. \square

We now establish that the given history is wait-free. The Writer procedure is obviously wait-free. However, it is not obvious that the `do` statement of a Read operation terminates (recall that L might be infinite). This is established below; the following four lemmas are used in the proof. The first of these lemmas is also used in the proof of Read Precedence.

Lemma 1: The following two assertions are invariants.

$$\begin{aligned} WP &= wp \vee RP = rp \\ WP &= rp \vee RP = d \end{aligned}$$

Proof: To prove that the given assertions are invariants, we consider the assertions $B0, \dots, B4$ defined below and show that $B0 \vee \dots \vee B4$ is an invariant.

$$B0 \equiv d \quad = wp = WP \quad = rp \quad = RP$$

$$\begin{aligned}
B1 &\equiv d \oplus 1 = wp = WP \oplus 1 = rp \oplus 1 = RP \oplus 1 \\
B2 &\equiv d \oplus 1 = wp = WP = rp \oplus 1 = RP \oplus 1 \\
B3 &\equiv d \oplus 1 = wp = WP = rp = RP \oplus 1 \\
B4 &\equiv d \oplus 1 = wp = WP = rp = RP
\end{aligned}$$

To see that $B0 \vee \dots \vee B4$ is an invariant, observe the following.

- B0 is initially true, and the only statement that can possibly falsify it is statement 2 of the Writer. But, executing this statement when B0 is true establishes B1.
- The only statement that can possibly falsify B1 is statement 7 of the Writer. But, executing this statement when B1 is true establishes B2.
- The only statement that can possibly falsify B2 is statement 0 of the Reader. But, executing this statement when B2 is true establishes B3.
- The only statement that can possibly falsify B3 is statement 1 of the Reader. But, executing this statement when B3 is true establishes B4.
- The only statement that can possibly falsify B4 is statement 1 of the Writer. But, executing this statement when B4 is true establishes B0.

Thus, we conclude that $B0 \vee \dots \vee B4$ is an invariant. Observe that $(B0 \vee \dots \vee B4) \Rightarrow (WP = wp \vee RP = rp)$, and $(B0 \vee \dots \vee B4) \Rightarrow (WP = rp \vee RP = d)$. This implies that both of the given assertions are invariants. \square

The next two lemmas show that no Write operation can modify $Y[i, j]$ while some Read operation is reading it.

Lemma 2: Let r be a Read operation and let w be a Write operation. Then, $after(w:2) \wedge \neg after(w:6) \wedge after(r:2) \wedge \neg after(r:5) \Rightarrow (rp \neq wp) \vee (ralt \neq walt)$.

Proof: Let r be a Read operation and let w be a Write operation. Define the assertions B and C as follows.

$$\begin{aligned}
B &\equiv after(w:2) \wedge \neg after(w:6) \wedge after(r:2) \wedge \neg after(r:5) \\
C &\equiv B \Rightarrow (ralt = Z[WP]) \vee (rp \neq wp) \vee (rp \neq WP)
\end{aligned}$$

We begin by proving that assertion C is an invariant. Note that B is initially false. This implies that C is initially true. Therefore, to prove that C is an invariant, we are obligated to show that C is never falsified. Note that if B is true at some state, then no event that is enabled at that state can change the value of $ralt$, $Z[0]$, $Z[1]$, WP , rp , or wp . Therefore, the only events that can possibly falsify C are those events that make B true. There are two such events, namely, $r:2$ and $w:2$.

First, consider the event $r:2$. By the text of the Reader procedure, $ralt = Z[rp]$ holds at the state following $r:2$. As we now show, $(ralt = Z[rp]) \Rightarrow C$.

$$(ralt = Z[rp]) = (ralt = Z[rp]) \wedge (rp = WP \vee rp \neq WP)$$

$$\begin{aligned}
&\Rightarrow (ralt = Z[WP]) \vee (rp \neq WP) \\
&\Rightarrow C
\end{aligned} \tag{1}$$

This implies that C holds at the state following $r:2$. Thus, we conclude that the event $r:2$ does not falsify C .

Now, consider the event $w:2$. Let t be the state prior to $w:2$, and let u be the state following $w:2$. Our proof obligation is to show that C holds at state u . If B is false at u , then C holds at u . So, assume that B holds at u . This implies that $u \models after(r:2)$. Because u is reached from t via the occurrence of the event $w:2$, we also have $t \models after(r:2)$. This implies that $r:2 \prec w:2$. By the text of the Reader procedure, $r:1 \prec r:2$. This establishes the following precedence assertion.

$$r:1 \prec r:2 \prec w:2 \tag{2}$$

Consider the two events $r:1$ and $w:1$. Either $r:1 \prec w:1$ or $w:1 \prec r:1$. We first dispose of the former case. By the text of the Writer procedure, $w:1 \prec w:2$. Therefore,

$$r:1 \prec w:1 \prec w:2 . \tag{3}$$

Because B holds at u , $after(r:5)$ is false at u . (Thus, if the event $r:5$ occurs in the given history, then by the definition of state u , we have $w:2 \prec r:5$. However, in the proof of wait-freedom, we cannot assume that the given history is well-formed. So, we cannot assert that the event $r:5$ occurs in the given history.) Because u is the state following the event $w:2$, by (3), this implies that $w!d$ is determined by $r:1$. Therefore, $w!d = r!rp$. By the text of the Writer procedure, $w!wp \neq w!d$. Hence, $w!wp \neq r!rp$. Because B holds at u , by the text of the Reader and Writer procedures, $u \models rp = r!rp \wedge wp = w!wp$. Therefore, $u \models rp \neq wp$. This implies that $u \models C$, which establishes our proof obligation.

We now consider the other case mentioned above, i.e., $w:1 \prec r:1$. In this case, by (2), we have the following precedence assertion.

$$w:1 \prec r:1 \prec r:2 \prec w:2$$

By the text of the Reader procedure, $ralt = Z[rp]$ holds at the state following $r:2$. Because $after(r:5)$ is false at u , and because u is the state following $w:2$, the above precedence assertion implies that $ralt$, rp , $Z[0]$, and $Z[1]$ each have the same value both at the state following $r:2$ and at state u . Therefore, $u \models ralt = Z[rp]$. By (1), this implies that $u \models C$. Thus, we conclude that the event $w:2$ does not falsify C . This completes the proof that C is an invariant.

We now use the fact that C is an invariant to establish our proof obligation. Consider a state t' such that B holds. Our proof obligation is to show that $t' \models (rp \neq wp) \vee (ralt \neq walt)$. By the definition of B , we have $t' \models after(r:2) \wedge \neg after(w:6)$. This implies that w does not precede r . Therefore, w is not the initial Write operation, because, by assumption, the initial Write operation precedes all Read operations. So, let v be the Write operation that immediately precedes w . Because B holds at state t' , we have $t' \models after(w:2) \wedge \neg after(w:6)$. Therefore, because v and w are consecutive Write operations, the following assertion holds.

$$t' \models Z[v!wp] = v!walt \wedge WP = v!wp \wedge walt = w!walt \tag{4}$$

Because v and w are consecutive Write operations, $v!walt \neq w!walt$. Therefore, by (4), we have the following.

$$t' \models Z[WP] \neq walt \quad (5)$$

Because B holds at t' , and because C is an invariant, we have $t' \models (ralt = Z[WP]) \vee (rp \neq wp) \vee (rp \neq WP)$. If $t' \models (ralt = Z[WP]) \vee (rp \neq wp)$, then by (5), $t' \models (ralt \neq walt) \vee (rp \neq wp)$. So, in this case our proof obligation is satisfied. Now, consider the other possibility, i.e., assume that $t' \models rp \neq WP$. By Lemma 1, $t' \models (WP = rp) \vee (RP = d)$. Therefore, because $t' \models rp \neq WP$, we have $t' \models RP = d$. Because B holds at t' , by the text of the Reader and Writer procedures, $t' \models rp = RP \wedge wp = d \oplus 1$. This implies that $t' \models rp \neq wp$, which establishes our proof obligation. \square

Lemma 3: Let r be a Read operation, and let c equal either 0 or 1. Then, $(\forall k :: Y[rp, ralt][k] = c \wedge after(r:2) \text{ unless } after(r:5))$.

Proof: Let r and c be as defined in the statement of the lemma, and let $k \geq 0$. Let t and u be consecutive states such that $t \models Y[rp, ralt][k] = c \wedge after(r:2)$ and $u \models \neg after(r:5)$. Because $after(r:2)$ holds at t and $after(r:5)$ is false at u , the following assertion holds at both states t and u .

$$rp = r!rp \wedge ralt = r!ralt$$

Let $i = r!rp$ and $j = r!ralt$. Then, $t \models Y[i, j][k] = c$. By the definition of *unless*, our proof obligation is to show that $u \models Y[i, j][k] = c$.

Suppose, to the contrary, that $u \models Y[i, j][k] \neq c$. Then, the event prior to state u is an event of some Write operation w , and this event writes to $Y[i, j][k]$. By the text of the Writer procedure, this implies that

$$t \models after(w:2) \wedge \neg after(w:6) \quad (6)$$

Assertion (6) implies that $t \models wp = w!wp \wedge walt = w!walt$. Because an event of w writes to $Y[i, j][k]$, $w!wp = i$ and $w!walt = j$. Therefore, $t \models wp = i \wedge walt = j$. Hence, by the definition of i and j ,

$$t \models wp = rp \wedge walt = ralt \quad (7)$$

Because t and u are consecutive states and $t \models after(r:2)$ and $u \models \neg after(r:5)$, we have

$$t \models after(r:2) \wedge \neg after(r:5) \quad (8)$$

By (6), (7), and (8) we have a contradiction of Lemma 2 at state t . Therefore, our assumption that $u \models Y[i, j][k] \neq c$ is false, i.e., $u \models Y[i, j][k] = c$. \square

The following invariant is used in the proof of wait-freedom.

Lemma 4: $(\forall i, j :: (\exists k :: Y[i, j][2k] = 0))$.

Proof: The given assertion is initially true by the definition of the initial state. It remains true because a Write operation assigns the value 1 to $Y[i, j][2k]$, for some k , only after it has assigned

the value 0 to some $Y[i, j][2n]$, where $n > k$. Therefore, the given assertion is an invariant. \square

Proof of Wait-Freedom: We prove that the given history is wait-free by showing that each Read operation consists of a finite number of events. Let r be a Read operation. If the event $r:2$ does not occur in the given history, then r is incomplete and consists of a finite number of events. So, assume that the event $r:2$ occurs in the given history. By Lemma 3, for each $k \geq 0$, $Y[rp, ralt][k]$ has the same value at each state in the interval of states for which $after(r:2) \wedge \neg after(r:5)$ holds. Therefore, by Lemma 4, there exists some k such that $Y[rp, ralt][2k] = 0$ at each state in this interval. This implies that the **do** statement of r does not iterate an infinite number of times. Thus, r consists of a finite number of events. \square

In this rest of this appendix, we assume that the given history is well-formed. (Because the given history is well-formed, we can assume that each operation considered hereafter is complete.) We show that the history is serializable by defining a function ϕ , and by proving that the defined ϕ satisfies the conditions of Uniqueness, Integrity, Proximity, and Read Precedence of the Shrinking Lemma. Note that the definition of ϕ depends on the auxiliary *id* variables.

Definition: Let r be a Read operation and let w be Write operation. Then, $\phi(r) \equiv r!id$ and $\phi(w) \equiv w!id$. \square

We now establish the conditions of the Shrinking Lemma. The following lemma is used in the proof. Note that this lemma is a strengthening of the Integrity condition.

Lemma 5: Let r be a Read operation and let $i = r!rp$. Then, there exists a Write operation w such that: (i) $last(Z[i]) = w$ at the state prior to $r:2$, (ii) $\phi(w) = \phi(r)$, and (iii) $w!val = r!val$.

Proof: Let r and i be as defined in the lemma, and let t be the state prior to the event $r:2$.

We first establish (i). By our assumption concerning the initial Writes, there exists a Write operation v such that $v:7$ determines $r!rp$. This implies that $v!wp = r!rp$; hence, by the definition of i , $v!wp = i$. Thus, event $v:6$, which occurs before state t , writes to $Z[i]$. This implies that there exists w such that $t \models last(Z[i]) = w$, establishing (i).

Because $t \models last(Z[i]) = w$, by the text of the Writer procedure, $w!wp = i$. Therefore, by the definition of i ,

$$w!wp = r!rp \quad . \quad (9)$$

Because $t \models last(Z[i]) = w$, we have $t \models Z[i] = w!walt$. Because the Writer updates $Q[i]$ and $Z[i]$ in a single action, $t \models last(Q[i]) = w$. Thus, $t \models Q[i] = w!id$. Because t is the state prior to $r:2$, $t \models Z[i] = r!ralt \wedge Q[i] = r!id$. Therefore, by transitivity, the following assertion holds.

$$w!walt = r!ralt \wedge w!id = r!id \quad (10)$$

By the definition of ϕ , (10) implies that $\phi(w) = \phi(r)$, which establishes (ii). We now meet our remaining proof obligation (iii), i.e., $w!val = r!val$. Let u be the state following $r:2$. Let $j = r!ralt$, and let $m = w!n$. By the text of the Reader procedure, $rp = i$ and $ralt = j$ at each state in the

interval between $r:2$ and $r:5$. Hence, by Lemma 3, each element of $Y[i, j]$ has the same value at each state in this interval. Therefore, to prove that $w!val = r!val$, it suffices to prove that the following expression holds.

$$u \models Y[i, j][2m] = 0 \wedge (\forall k : 0 \leq k < m : Y[i, j][2k] = 1 \wedge Y[i, j][2k + 1] = w!val[k])$$

By the text of the Writer procedure, it is sufficient to show that $u \models (\forall k : 0 \leq k \leq 2m : last(Y[i, j][k]) = w)$. So, let $0 \leq k \leq 2m$. As stated previously, $w!wp = i$. By (10) and the definition of j , $w!walt = j$. Hence, by the definition of k , there exists an event e of Write operation w that writes to $Y[i, j][k]$. By the text of the Writer procedure, $e \prec w:6$. Because $t \models last(Z[i]) = w$, $w:6 \prec r:2$. Therefore, $e \prec w:6 \prec r:2$. This implies that either $u \models last(Y[i, j][k]) = w$ or $u \models last(Y[i, j][k]) = w'$, where w' is a successor of w .

Thus, to establish (iii), it suffices to prove that if w' is a successor of w , then $u \models last(Y[i, j][k]) \neq w'$. So, suppose that w' is a successor of w . If $w'!wp \neq i$ or $w'!walt \neq j$, then clearly $u \models last(Y[i, j][k]) \neq w'$. So, assume that $w'!wp = i$ and $w'!walt = j$. In this case, we prove that $u \models last(Y[i, j][k]) \neq w'$ by showing that $r:2 \prec w':2$. Assume, to the contrary, that $w':2 \prec r:2$. By the definition of i and j ,

$$w'!wp = r!rp \wedge w'!walt = r!ralt \quad . \quad (11)$$

Because $t \models last(Z[i]) = w$ and w precedes w' , $r:2 \prec w':6$. Therefore, $w':2 \prec r:2 \prec w':6$. This precedence assertion implies that $u \models wp = w'!wp \wedge walt = w'!walt \wedge rp = r!rp \wedge ralt = r!ralt$. Therefore, by (11), $u \models wp = rp \wedge walt = ralt$. This implies that Lemma 2 is violated at state u . Thus, our assumption that $w':2 \prec r:2$ is false, and we conclude that $r:2 \prec w':2$. This establishes (iii). \square

Proof of Uniqueness: Uniqueness is satisfied because each Write operation increments the private variable id of the Writer. \square

Proof of Integrity: Integrity follows from conditions (ii) and (iii) of Lemma 5. \square

Proof of Proximity: Let r be a Read operation and let v be a Write operation. Let i and w be as defined in Lemma 5, and let t be the state prior to the event $r:2$. Then, by condition (ii) of the lemma, $\phi(r) = \phi(w)$. Our proof obligation is to show that if r precedes v , then $\phi(r) < \phi(v)$, and if v precedes r , then $\phi(v) \leq \phi(r)$. We consider the two cases separately.

Case 1: r precedes v . By condition (i) of Lemma 5, $t \models last(Z[i]) = w$. Therefore, $w:6 \prec r:2$. Because r precedes v (and because the Write operations are totally ordered), this implies that w precedes v . Hence, by Uniqueness, $\phi(w) < \phi(v)$. Therefore, because $\phi(r) = \phi(w)$, we have $\phi(r) < \phi(v)$.

Case 2: v precedes r . Because $\phi(r) = \phi(w)$, to prove that $\phi(v) \leq \phi(r)$, it suffices to prove that $\phi(v) \leq \phi(w)$. If $v = w$, then our proof obligation is satisfied. So, assume that $v \neq w$. By the Uniqueness condition, it suffices to prove that v precedes w .

Let w' be the Write operation such that $w':7$ determines $r!rp$. (Note that w' exists by our assumption concerning the initial Writes.) Then, by the text of the Reader and Writer procedures, w' precedes r and $w'!wp = r!rp$. Therefore, by the definition of i , $w'!wp = i$. Hence, $w':6$ writes to $Z[i]$. Because $w':7$ determines $r!rp$ and v precedes r , $v:7 \preceq w':7$. Because the Write operations are totally ordered, this implies that v precedes or equals w' .

Because the Write operations are totally ordered, either v precedes w or w precedes v . In the latter case, because v precedes or equals w' , w precedes w' . Because w' precedes r and $w':6$ writes to $Z[i]$, this implies that $t \models \text{last}(Z[i]) \neq w$. However, this contradicts condition (i) of Lemma 5. Thus, we conclude that v precedes w . \square

Proof of Read Precedence: Let r and s be two Read operations, and assume that r precedes s . Our proof obligation is to show that $\phi(r) \leq \phi(s)$. By condition (ii) of Lemma 5, there exist Write operations v and w such that $\phi(v) = \phi(r)$ and $\phi(w) = \phi(s)$. Thus, it suffices to prove that $\phi(v) \leq \phi(w)$. Assume, to the contrary, that $\phi(w) < \phi(v)$. Then, by Uniqueness (and by the fact that the Write operations are totally ordered), w precedes v .

Let $i = r!rp$ and $j = s!rp$. By the text of the Writer procedure, $w:0 \prec w:7$. Because w precedes v , $w:7 \prec v:0$. By the text of the Writer procedure, $v:0 \prec v:6$. By condition (i) of Lemma 5, $\text{last}(Z[i]) = v$ at the state prior to $r:2$; therefore, $v:6 \prec r:2$. By the text of the Reader procedure, $r:2 \prec r:5$. Because r precedes s , $r:5 \prec s:0$. This establishes the following precedence assertion.

$$w:0 \prec w:7 \prec v:0 \prec v:6 \prec r:2 \prec r:5 \prec s:0$$

Because $\phi(s) = \phi(w)$, by Uniqueness and Proximity, no Write operation succeeds w and precedes s . Hence, by the above precedence assertion, w and v are consecutive Write operations and v does not precede s . This implies that $s:0 \prec v:7$. Therefore, we have the following assertion.

$$w:0 \prec w:7 \prec v:0 \prec v:6 \prec r:2 \prec r:5 \prec s:0 \prec v:7 \tag{12}$$

We now show that $r!rp = s!rp = v!wp = w!wp$. Note that condition (i) of Lemma 5 implies that $v!wp = i$ and $w!wp = j$. Thus, by the definition of i and j , we have the following assertion.

$$r!rp = v!wp \wedge s!rp = w!wp \tag{13}$$

Let s' be the Read operation that is the immediate successor of r . (Note that s' exists because r precedes s .) Then, $r:5 \prec s':0$. Moreover, because r precedes s , s' precedes or equals s ; hence, $s':0 \preceq s:0$. Therefore, by (12), the following precedence assertion holds.

$$w:0 \prec w:7 \prec v:0 \prec v:6 \prec r:2 \prec r:5 \prec s':0 \prec v:7 \tag{14}$$

Because w and v are consecutive Write operations, this precedence assertion implies that $s'!rp$ is determined by $w:7$. Thus,

$$s'!rp = w!wp \ . \tag{15}$$

Let t be the state following the event $s':0$. Because w and v are consecutive Write operations and r and s' are consecutive Read operations, by (14),

$$t \models WP = w!wp \wedge wp = v!wp \wedge RP = r!rp \wedge rp = s'!rp \ .$$

By (15) and (13), $s!rp = s!rp$. Also, by Lemma 1, $t \models WP = wp \vee RP = rp$. Therefore, $w!wp = v!wp \vee r!rp = s!rp$. Consequently, by (13),

$$r!rp = s!rp = v!wp = w!wp .$$

Because $v!wp = w!wp = j$ and w and v are consecutive, $v:6$ overwrites the value written to $Z[j]$ by $w:6$. By assertion (12) this implies that $last(Z[j]) \neq w$ at the state prior to $s:2$. However, by condition (i) of Lemma 5, $last(Z[j]) = w$ at this state. Therefore, we have a contradiction. Thus, our assumption that $\phi(w) < \phi(v)$ is false. \square

References

- [1] Y. Afek, H. Attiya, D. Dolev, E. Gafni, M. Merritt, and N. Shavit, "Atomic Snapshots," *Proceedings of the Ninth Annual Symposium on Principles of Distributed Computing*, 1990, to appear.
- [2] J. Anderson, "Composite Registers," *Proceedings of the Ninth Annual Symposium on Principles of Distributed Computing*, 1990, to appear. Also available as Technical Report TR.89.25, Department of Computer Sciences, University of Texas at Austin, September 1989.
- [3] J. Anderson, "Multiple-writer Composite Registers," Technical Report TR.89.26, Department of Computer Sciences, University of Texas at Austin, September 1989.
- [4] J. Anderson, *Atomicity of Concurrent Programs*, Ph.D. Dissertation, Department of Computer Sciences, University of Texas at Austin, 1990.
- [5] J. Anderson and M. Gouda, "The Virtue of Patience: Concurrent Programming With and Without Waiting," unpublished manuscript, Department of Computer Sciences, University of Texas at Austin, 1988.
- [6] B. Bloom, "Constructing Two-Writer Atomic Registers," *IEEE Transactions on Computers*, Vol. 37, No. 12, December 1988, pp. 1506-1514.
- [7] J. Burns and G. Peterson, "Constructing Multi-Reader Atomic Values from Non-Atomic Values," *Proceedings of the Sixth Annual Symposium on Principles of Distributed Computing*, 1987, pp. 222-231.
- [8] J. Burns and G. Peterson, "Pure Buffers for Concurrent Reading While Writing," Technical Report GIT-ICS-87/17, School of Information and Computer Science, Georgia Institute of Technology, April 1987.
- [9] B. Chor, A. Israeli, and M. Li, "On Processor Coordination Using Asynchronous Hardware," *Principles of the Sixth Annual Symposium on Principles of Distributed Computing*, 1987, pp. 86-97.
- [10] E. Dijkstra, *A Discipline of Programming*, Prentice-Hall, Englewood Cliffs, New Jersey, 1976.

- [11] M. Gouda, "Serializable Programs, Parallelizable Assertions: A Basis for Interleaving," in *Beauty is Our Business*, eds. W.H.J. Feijen, A.J.M. van Gasteren, D. Gries, and J. Misra, Springer-Verlag, 1990, pp. 135-140.
- [12] M. Herlihy, "Wait-Free Implementation of Concurrent Objects," *Proceedings of the Seventh Annual Symposium on Principles of Distributed Computing*, 1988.
- [13] M. Herlihy and J. Wing, "Axioms for Concurrent Objects," *Proceedings of the 4th Annual ACM Symposium on Principles of Programming Languages*, 1987.
- [14] C.A.R. Hoare, "Towards a Theory of Parallel Programming," *Operating Systems Techniques*, eds. Hoare and Perott, Academic Press, New York, 1972.
- [15] L. Kirousis, E. Kranakis, and P. Vitanyi, "Atomic Multireader Register," Proceedings of the Second International Workshop on Distributed Computing, Springer Verlag Lecture Notes in Computer Science 312, 1987, pp. 278-296.
- [16] L. Lamport, "On Interprocess Communication, Parts I and II," *Distributed Computing*, Vol. 1, 1986, pp. 77-101.
- [17] L. Lamport, "*win* and *sin*: Predicate Transformers for Concurrency," Technical Report, Systems Research Center, Digital Equipment Corporation, October 1986.
- [18] L. Lamport and F. Schneider, "The Hoare Logic of CSP, and All That," *ACM Transactions on Programming Languages and Systems*, Vol. 6, No. 2, April 1984, pp. 281-296.
- [19] M. Li, J. Tromp, and P. Vitanyi, "How to Construct Wait-Free Variables," unpublished manuscript, Centrum voor Wiskunde en Informatica, Amsterdam, The Netherlands, 1989.
- [20] M. Merritt, private communication, 1990.
- [21] Z. Manna and A. Pnueli, "Adequate Proof Principles for Invariance and Liveness Properties of Concurrent Programs," *Science of Computer Programming*, Vol. 4, 1984, pp. 257-289.
- [22] J. Misra, "Axioms for Memory Access in Asynchronous Hardware Systems," *ACM Transactions on Programming Languages and Systems*, Vol. 8, No. 1, January 1986, pp. 142-153.
- [23] R. Newman-Wolfe, "A Protocol for Wait-Free, Atomic, Multi-Reader Shared Variables," *Proceedings of the Sixth Annual Symposium on Principles of Distributed Computing*, 1987, pp. 232-248.
- [24] S. Owicki and D. Gries, "An Axiomatic Proof Technique for Parallel Programs I," *Acta Informatica*, Vol. 6, 1976, pp. 319-340.
- [25] G. Peterson and J. Burns, "Concurrent Reading While Writing II: The Multi-Writer case," *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*, 1987.
- [26] A. Singh, J. Anderson, and M. Gouda, "The Elusive Atomic Register, Revisited," *Proceedings of the Sixth Annual Symposium on Principles of Distributed Computing*, 1987, pp. 206-221.

- [27] J. Tromp, "How to Construct an Atomic Variable," unpublished manuscript, Centrum voor Wiskunde en Informatica, Amsterdam, The Netherlands, 1989.
- [28] P. Vitanyi and B. Awerbuch, "Atomic Shared Register Access by Asynchronous Hardware," *Proceedings of the 27th IEEE Symposium on the Foundations of Computer Science*, 1986, pp. 233-243.