

Fail-Stop Processors: An Approach to Designing Fault-Tolerant Computing Systems

RICHARD D. SCHLICHTING

University of Arizona

and

FRED B. SCHNEIDER

Cornell University

A methodology that facilitates the design of fault-tolerant computing systems is presented. It is based on the notion of a fail-stop processor. Such a processor automatically halts in response to any internal failure and does so before the effects of that failure become visible. The problem of implementing processors that, with high probability, behave like fail-stop processors is addressed. Axiomatic program verification techniques are described for use in developing provably correct programs for fail-stop processors. The design of a process control system illustrates the use of our methodology.

Categories and Subject Descriptors: C.2.4 [Computer-Communications Networks]: Distributed Systems—network operating systems; C.3 [Special-Purpose and Application-Based Systems]:—real-time systems; D.4.5 [Operating Systems]: Reliability—verification; F.3.1 [Logics and Meaning of Programs]: Specifying and Verifying and Reasoning about Programs

General Terms: Reliability, Verification

Additional Key Words and Phrases: Fail-stop

1. INTRODUCTION

Programming a computer system that is subject to failures is a difficult task. A malfunctioning processor might perform arbitrary and spontaneous state transformations instead of the transformations specified by the programs it executes. Thus even a correct program cannot be counted on to implement a desired input-output relation when executed on a malfunctioning processor. On the other hand, it is impossible to build a computer system that always operates correctly in spite

This work is supported in part by NSF Grant MCS-81-03605.

Authors' addresses: R. D. Schlichting, Department of Computer Science, University of Arizona, Tucson, Arizona 85721; F. B. Schneider, Department of Computer Science, Cornell University, Ithaca, New York 14853.

This paper was originally submitted to *ACM Transactions on Programming Languages and Systems*. The responsible editor was Susan L. Graham. The authors and editor kindly agreed to transfer the paper to the *ACM Transactions on Computer Systems*.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1983 ACM 0734-2071/83/0800-0222 \$00.75

ACM Transactions on Computing Systems, Vol. 1, No. 3, August 1983, Pages 222-238.

of failures in its components by using (only) a finite amount of hardware.¹ Thus the goal of implementing completely fault-tolerant computing systems is unattainable. Fortunately, most applications do not require complete fault-tolerance. Rather, it is sufficient that a system work correctly provided that no more than some predefined number of failures occur within some time interval, or that certain types of failures do not occur. This more modest goal *is* attainable.

In this paper we present an approach to designing fault-tolerant computing systems based on the notion of a *fail-stop processor*, a processor with well-defined failure-mode operating characteristics. Briefly, our approach is as follows. First, software is designed assuming the existence of a computing system composed of one or more fail-stop processors; the number of processors required is dictated by response-time constraints that must be satisfied by the system. Then, a computing system is designed that implements the requisite fail-stop processors.

We proceed as follows. Section 2 describes the characteristics of a fail-stop processor and considers how such processors can be approximated using present-day hardware. Section 3 discusses extensions to axiomatic verification techniques for facilitating the development of provably correct programs for fail-stop processors. Satisfying response time constraints in the presence of failures is the subject of Section 4. Section 5 discusses the application of our approach to a nontrivial problem: the design of a fault-tolerant process-control system. Section 6 contrasts our work with other approaches to designing fault-tolerant systems, and Section 7 presents some conclusions.

2. FAIL-STOP PROCESSORS

2.1 Definition

A processor is characterized by its instruction set. Each instruction causes a well-defined transformation on the internal state of the processor and/or the connected storage and peripheral devices. Thus the effects of executing each instruction can be described by a precise semantic definition, be it a temporal axiomatization of the instruction set [18] or a "Principles of Operation" manual. A *failure* occurs when the behavior of the processor is not consistent with this semantic definition.

A *fail-stop processor* is distinguished by its extremely simple failure-mode operating characteristics. First, the internal state of a fail-stop processor and some predefined portion of the connected storage are assumed to be *volatile*; the contents of volatile storage are irretrievably lost whenever a failure occurs. The remaining storage is defined to be *stable*; it is unaffected by any kind of failure. Secondly, in contrast to a real processor, a fail-stop processor never performs an erroneous state transformation due to a failure. Instead, the processor simply halts. Thus, the only visible effects of a failure in a fail-stop processor are

FS1: It stops executing.

FS2: The internal state and contents of the volatile storage connected to it are lost.

¹ *Sed quis custodiet ipsos Custodes?* (Who shall guard the guards themselves?), Juvenal, Satires VI, 347.

2.2 Implementation

While the notion of a fail-stop processor is a useful abstraction, it is impossible to implement using a finite amount of hardware. With only a finite amount of hardware, a sufficient number of failures could disable all the error detection mechanisms and thus allow arbitrary behavior. It is, however, possible to construct computing systems that, with high probability, approximate the behavior of a fail-stop processor.

One approach is to construct a system that behaves as specified, unless too many failures occur within some specified time interval, after which no assumptions about its behavior can be made. A *k*-fail-stop processor is a computing system that behaves like a fail-stop processor unless $k + 1$ or more failures occur in its components. One implementation of a *k*-fail-stop processor is described below; another appears in [23].

A *k*-fail-stop processor can be implemented by exploiting any solution to the Byzantine Generals Problem [13] (or the Interactive Consistency Problem [17]). Such a protocol allows a collection of processors to agree on a value sent by a potentially faulty transmitter (“Commanding General”), so that

- (1) each nonfaulty processor agrees on the value sent by the transmitter, and
- (2) if the transmitter is nonfaulty, each nonfaulty processor receives the value sent by the transmitter.

A number of real processors and volatile memory units are interconnected by a communications network to form a single *k*-fail-stop processor and its attendant stable storage. Each memory unit M_i is read by all processors but written to by only one, P_i . Failures are detected by having each processor run the same program and comparing results. Thus a copy of each variable is stored in each memory unit. During execution, whenever the value of a variable from stable storage is required, the value of that variable is read from each memory unit and a solution to the Byzantine Generals Problem is employed to distribute the vector of values read to every processor. If all of the values are not identical, then a failure has occurred and it is signaled. (Nonfaulty processors will halt when the failure is signaled.) A total of $2k + 1$ processors are required in order for up to k failures to be tolerated.

Since processors execute asynchronously, execution of the replicated programs must be synchronized to compare results. This can be accomplished, assuming processor clocks run at roughly the same rate, by associating a logical clock [11] with each program. This logical clock is incremented whenever a variable that is supposed to be stored in the stable storage of the *k*-fail-stop processor is read or written. To synchronize, a processor constructs a vector of the values of each of these clocks, again using a solution to the Byzantine Generals Problem, and marks time until all components in the vector have the same value or a “time-out” period has elapsed. In the latter case, a failure has occurred and it is signaled.

A collection of *k*-fail-stop processors are interconnected to implement a system in which real-time response constraints can be met despite failures. In that case, it must be possible for one *k*-fail-stop processor to detect that another has stopped and then to read the contents of that *k*-fail-stop processor’s stable storage so that the computation in progress at the time of the failure can be continued. This is

accomplished as follows. Each k -fail-stop processor is connected to a communications network that allows it to read the contents of the memory units that make up the other k -fail-stop processors. A special location, *failed_i*, in each memory unit M_i , is reserved to record whether processor P_i thinks that the fail-stop processor it is a part of has halted due to a failure. A k -fail-stop processor fsp determines that another, fsp' has halted by computing the vector of values *failed_i* for each memory unit in fsp' , using a solution to the Byzantine Generals Problem. If any of the components has value true, then fsp' is deemed halted. Should fsp require the values of the variables in the stable storage of fsp' , they can be reconstructed as follows. Each processor in fsp reads a different one of the memory units that make up fsp' . Then, using a solution to the Byzantine Generals Problem, these values are exchanged. The majority value is taken to be the value of the variable. Since at most k of the values read from fsp' are wrong, at least $2k + 1$ different memory units are required to implement stable storage.

While the feasibility of implementing fail-stop processors is established by this argument, the practicality is not. However, recent work in the implementation of highly reliable processors gives reason to believe that it is indeed practical to implement fail-stop processor approximations. Both FTMP [10] and SIFT [26] could be configured to behave like a collection of fail-stop processor approximations; both employ replicated processor and memory units. Redundancy can also be introduced at lower levels in a variety of ways [1, 25]. The level at which redundancy is applied is an important issue and is treated in [2].

3. PROGRAMMING A FAIL-STOP PROCESSOR

3.1 Recovery Protocols

A program executing on a fail-stop processor is halted when a failure occurs. Execution may then be *restarted* on a correctly functioning fail-stop processor. (This may be the original processor if the cause of the failure has been repaired, or it may be another fail-stop processor.) When a program is restarted, the internal processor state and the contents of volatile storage are unavailable. Thus, some routine is needed that can complete the state transformation that was in progress at the time of the failure and restore storage to a well-defined state. Such a routine is called a *recovery protocol*.

Clearly, a recovery protocol (1) must execute correctly when started in any intermediate state that could be visible after a failure and (2) can only use information that is in stable storage. In addition, because the code for a recovery protocol must be available after a failure, it must be kept in stable storage.

We associate a recovery protocol R with a sequence of statements A called the *action statement* to form a *fault-tolerant action FTA* as follows:

```
FTA: action
      A
      recovery
      R
      end
```

Execution of FTA consists of establishing R as the recovery protocol to be in effect when A is executed and then executing A . If execution of FTA is interrupted

by a failure, upon restart execution continues with the recovery protocol in effect. Subsequent failures cause execution of FTA to be halted and execution of the recovery protocol in effect to begin anew when the program is restarted. Execution of FTA terminates when execution of either A or R is performed in its entirety without interruption. At that time, either the recovery protocol in effect when FTA started is reestablished, or, if another fault-tolerant action FTA' follows FTA, then the recovery protocol for FTA' is established.

The following syntactic abbreviation is used to denote that an action statement A serves as its own recovery protocol:

FTA: **action, recovery**
 A
end

Such a fault-tolerant action is called a *restartable action*.²

A program running on a fail-stop processor must at all times have a recovery protocol in effect. This will be the case if the program itself is a single fault-tolerant action. Alternatively, a program can be structured as a sequence of fault-tolerant actions, assuming that establishment of a recovery protocol can be done in such a way that at all times either the old recovery protocol or the new one is in effect. Such an assumption seems quite reasonable.

3.2 Axioms for Fault-Tolerant Actions

Following the Floyd-Hoare axiomatic approach [8], an *assertion* is a Boolean-valued expression involving program and logical variables. The syntactic object,

$$\{P\} S \{Q\}$$

where P and Q are assertions and S is a programming language statement, is called a *triple*. The triple $\{P\} S \{Q\}$ is a *theorem* if there exists a proof of it in a specified formal deductive system, usually called a *programming logic*. A programming logic consists of a set of axioms and rules of inference that relate assertions, programming language statements, and triples. Of particular interest are those logics that are sound with respect to execution of programming language statements on the program state—that is, deductive systems that are consistent with the operation of a “real” machine. Then, the notation $\{P\} S \{Q\}$ is usually taken to mean:

If execution of S begins in a state in which P is true, and terminates, then Q will be true in the resulting state.

Numerous programming languages have been defined using such logics; a PASCAL-like language [9] extended with guarded commands [4] is used in this paper.

It is often more convenient to write a *proof outline* than a formal proof. A proof outline is a sequence of programming language statements interleaved with assertions. Each statement S in a proof outline is preceded directly by one assertion, called its *precondition* and denoted $pre(S)$, and is directly followed by an assertion, called its *postcondition* and denoted $post(S)$. A proof outline is an

² As we shall see, any fault-tolerant action can be converted to such a restartable action simply by omitting the action statement.

abbreviation for a proof if

PO1: for every statement S , the triple $\{\text{pre}(S)\} S \{\text{post}(S)\}$ is a theorem in the programming logic, and

PO2: whenever $\{P\}$ and $\{Q\}$ are adjacent in the proof outline, Q is provable from P .

Let FTA be a fault-tolerant action formed from action statement A and recovery protocol R . We wish to develop an inference rule that will allow derivation of

$$\{P\} \text{FTA} \{Q\}$$

as a theorem, while preserving the soundness of our programming logic with respect to execution on a fail-stop processor.

First, assume

F1: $\{P'\} A \{Q'\}$ and $\{P''\} R \{Q''\}$

have been proved. Then, for execution of A to establish Q , we will need

F2: $P \Rightarrow P'$ and $Q' \Rightarrow Q$.

Similarly, for the recovery protocol R to establish Q , the following (at least) must hold:

F3: $Q'' \Rightarrow Q$.

Recall that R is invoked only following a failure. By definition, the contents of volatile storage are undefined at that time. Therefore, any program variables needed for execution of R must be in stable storage.³ Thus, we require

F4: All program variables named in P'' must be in stable storage.

We must also ensure that whenever the recovery protocol receives control, stable storage is in a state that satisfies P'' . This will be facilitated by constructing a *replete proof outline*, a proof outline that contains assertions describing only those states that could be visible after a failure. Then, we will require that the precondition of the recovery protocol be satisfied in those states.

A *replete proof outline* is a proof outline in which certain assertions have been deleted so that

RPO1: No assertion appears between adjacent fault-tolerant actions.

RPO2: Every triple $\{P\} S \{Q\}$ in the replete proof outline satisfies either

- (a) S is a sequence of fault-tolerant actions, or
- (b) $\{P \vee Q\}$ is invariant over execution of S .

RPO1 and RPO2(a) follow because the program state that exists between the execution of two fault-tolerant actions FTA_1 and FTA_2 is never visible to the recovery protocol for the enclosing fault-tolerant action—either the recovery

³ If P'' is stronger than $\text{wp}(R, Q'')$, then variables may appear in P'' that need not be stored in stable storage. Thus, in the interest of minimizing the amount of stable storage used, the proofs should be in terms of the weakest assertions possible.

protocol for FTA_1 or the recovery protocol for FTA_2 will receive control. RPO2(b) follows because if $P \vee Q$ remains true while S is being executed, then either P or Q will be true of the state visible to the recovery protocol should a failure occur and both $\{P\}$ and $\{Q\}$ already appear as assertions in the replete proof outline.

For example, if

$$\{P\} \text{FTA}_1 \{P_1\} \text{FTA}_2 \{P_2\} \dots \text{FTA}_n \{P_n\}$$

is a proof outline, then

$$\{P\} \text{FTA}_1; \text{FTA}_2; \dots \text{FTA}_n \{P_n\}$$

is a replete proof outline. If assignment of an integer value to a variable is performed by executing a single, indivisible (store) instruction—as it is on most machines—then

$$\{x = 3\} x := 6 \{x = 6\}$$

is also a replete proof outline. This is because either the precondition or the postcondition of “ $x := 6$ ” is true of every state that occurs during execution of the assignment. Even if assignment is not implemented by execution of a single instruction,

$$\{val = 3\} x := val \{x = 3 \wedge val = 3\}$$

is a replete proof outline, because the assertion $\{val = 3\}$ is not destroyed by assignment to x ; it is true before, during, and after execution of “ $x := val$ ”.

Therefore, in addition to F1–F4, correct operation of a recovery protocol requires

F5: Given a fault-tolerant action with action statement A and recovery protocol R satisfying F1, let a_1, a_2, \dots, a_n be the assertions that appear in a replete proof outline of $\{P'\} A \{Q'\}$, and r_1, r_2, \dots, r_m be the assertions that appear in a replete proof outline of $\{P''\} R \{Q''\}$. Then:

- (a) $(\forall i: 1 \leq i \leq n: a_i \Rightarrow P'')$
- (b) $(\forall i: 1 \leq i \leq m: r_i \Rightarrow P')$

Lastly, it must be guaranteed that failures at processors other than the one executing FTA do not interfere with (i.e., invalidate) assertions in the proof outline of FTA. Suppose an assertion in FTA names variables stored in the volatile storage of another processor.⁴ Then, should that processor fail, such an assertion would no longer be true since the contents of volatile storage would have been lost. Hence, we require that

F6: Variables stored in volatile storage may not be named in assertions appearing in programs executing on other processors.

Given a fault-tolerant action, a restartable action that implements the same state transformation can always be constructed from the recovery protocol alone. (The proof of this follows from F3 and F5.) Thus, in theory, the action statement

⁴ This is often necessary when the actions of concurrently executing processes are synchronized. For example, if it is necessary to assert that a collection of processes are all executing in the same “phase” at the same time, then each would include assertions about the state of the others. See [21] for an example of such reasoning.

is unnecessary. In practice, the additional flexibility that results from having an action statement different from the recovery protocol is quite helpful. Presumably, failures are infrequent enough so that a recovery protocol can do a considerable amount of extra work in order to minimize the amount of (expensive) stable storage used. Use of such algorithms for normal processing would be unacceptable.

3.3 Fault-Tolerant Programs—A Simple Example

In addition to allowing axiomatic verification of programs written in terms of fault-tolerant actions, F1–F6 permit a programmer to develop a fault-tolerant program and its proof hand-in-hand, with the proof leading the way, as advocated in [4] and [6]. F4 allows those variables that must be stored in stable storage to be identified in a mechanical way from the proof; construction of a replete proof outline provides a mechanical way to determine the intermediate states that could be visible following a failure. To illustrate the use of rules F1–F6 as an aid in developing a recovery protocol, we consider the following (artificial) problem. (A more substantial example is treated in Section 5.)

Periodically, variables x and y are updated based on their previous values. Thus, given a function G , a routine called `update` is desired that runs on a fail-stop processor and satisfies the following specification:

$$\{P: x = X \wedge y = Y\} \text{ update } \{Q: x = G(X) \wedge y = G(Y)\}.$$

Logical variables X and Y represent the initial values of x and y , respectively.

If the possibility of failure is ignored, the following program will suffice:

```
S1: {P: x = X ∧ y = Y}
    S1a: x := G(x); {P1a: x = G(X) ∧ y = Y}
    S1b: y := G(y); {P1b: x = G(X) ∧ y = G(Y)}
    {Q: x = G(X) ∧ y = G(Y)}
```

Note that this is a replete proof outline, provided assignment is implemented as an atomic operation: $\{P \vee P1a\}$ is invariant over execution of $S1a$ and $\{P1a \vee P1b\}$ is invariant over execution of $S1b$.

Things become more complicated when the possibility of failure is considered. In particular, $S1$ could not be the action statement of a restartable action because F5 is violated (assuming G is not the identity function): both $P1a \Rightarrow P$ and $P1b \Rightarrow P$ are false. In order to construct a restartable action, we must find a way to make progress—compute $G(X)$ and $G(Y)$ —without destroying the initial values of x and y until both values have been updated. One way to do this is to modify $S1$ so that the new values are computed and stored in some temporary variables, giving the following restartable action:

```
U1: action, recovery
    {P: x = X ∧ y = Y}
    U1a: xnew := G(x); {x = X ∧ xnew = G(X) ∧ y = Y}
    U1b: ynew := G(y); {x = X ∧ xnew = G(X) ∧ y = Y ∧ ynew = G(Y)}
    end
    {Q': x = X ∧ xnew = G(X) ∧ y = Y ∧ ynew = G(Y)}
```


Note that in order to satisfy F4, x and y must be stored in stable storage but variables used in computing G need not be. Having established Q' , it is a simple matter to establish Q :

S2: $\{Q'' : xnew = G(X) \wedge ynew = G(Y)\}$
 S2a: $x := xnew; \{x = xnew = G(X) \wedge ynew = G(Y)\}$
 S2b: $y := ynew; \{x = xnew = G(X) \wedge y = ynew = G(Y)\}$
 $\{Q: x = G(X) \wedge y = G(Y)\}$

This is a replete proof outline, and provided $xnew$ and $ynew$ are stored in stable storage. F1-F6 are satisfied. So

U2: **action, recovery**
 $\{Q'' : xnew = G(X) \wedge ynew = G(Y)\}$
 U2a: $x := xnew; \{x = xnew = G(X) \wedge ynew = G(Y)\}$
 U2b: $y := ynew; \{x = xnew = G(X) \wedge y = ynew = G(Y)\}$
end
 $\{Q: x = G(X) \wedge y = G(Y)\}$

is a restartable action. Since $Q' \Rightarrow Q''$, the desired program is

U1; U2

4. TERMINATION AND RESPONSE TIME

Most statements in our programming notation are guaranteed to terminate, once started. However, loops and fault-tolerant actions are not. Techniques based on the use of variant functions or well-founded sets can be used for proving that a loop will terminate [4]. Unfortunately, without knowledge about the frequency of failures and statement execution times, termination of a program written in terms of fault-tolerant actions cannot be proved. This is because if failures occur with sufficiently high frequency, then there is no guarantee that the component fault-tolerant actions will terminate; neither the action statement nor the recovery protocol of a fault-tolerant action can be guaranteed to run without interruption, and so the recovery protocol could continually restart.

Moreover, such *liveness properties* [16] cannot even be expressed in a Hoare-style programming logic, like the one above. Thus, we must resort to informal means to argue that a program will terminate in a timely manner. Presumably, at some point in the future it will be possible to formalize such arguments. Harter and Bernstein [7] describe extensions to temporal logic [16] that allow construction of a proof that a program will meet some specific response-time goals. That work would have to be extended to deal with stochastically defined events for use in this context.

For a given execution of a program S on a fault-free processor, let $t(s)$ be the maximum length of time that elapses once execution of statement s is begun until execution of the next fault-tolerant action in S is started. Define

$$T_{\max} = \max_{s \in S} t(s).$$

For an execution of S to terminate at all, it is sufficient that there be (enough) intervals of length T_{\max} during which there are no failures. Then, no fault-tolerant action will be forever restarted as a result of the (high) frequency of failures.

Of course, this gives no bound on how much time will elapse before S completes. Rather, we have argued that S is guaranteed to terminate if the elapsed time between successive failures is long enough, often enough. This should not be surprising. However, it does provide some insight into how to structure a program in terms of fault-tolerant actions if frequent failures are expected: one should endeavor to minimize T_{\max} . This can be achieved by making entry into a fault-tolerant action a frequent event, either by nesting fault-tolerant actions, or by composing them in sequence.

Given a collection of fail-stop processors, it is possible to configure a system that not only implements a given relation between input and output, but performs this state transformation in a timely manner despite the occurrence of failures. After the failure of a fail-stop processor fsp , a *reconfiguration rule* is used to assign programs that were running on fsp to working fail-stop processors. The recovery protocol in effect at the time of the failure facilitates restart of the program. Thus, processor failures are transparent except for possibly increased execution times.

As a result of a failure, execution delays from the following sources are incurred:

- (1) Some time t_{detect} will elapse after the fail-stop processor halts until that fact is detected and reconfiguration is begun.
- (2) Reconfiguration causes execution delays, as well. First, t_{recon} is required to determine an appropriate assignment of programs to the remaining fail-stop processors. Then, t_{move} might be required to move the program code and contents of its stable storage.
- (3) In the worst case, the effects of the last T_A seconds execution by action statement A before the first failure will be lost.
- (4) T_R seconds worth of execution by repeated attempts to perform recovery protocol R as a result of each subsequent failure will also be lost.

Both T_A and T_R are defined for the specific execution that was interrupted.

This suggests the following strategy for constructing fault-tolerant systems that will continue to behave correctly in spite of up to k failures, for $k > 0$. A program is developed (1) that implements the desired state transformations when run on fail-stop processors, (2) that satisfies its real-time response constraints provided no failures occur, and (3) in which no process must respond to an event in less than T_F seconds, where

$$T_F = k(t_{\text{detect}} + t_{\text{recon}} + t_{\text{move}}) + (k - 1)T_R + T_A.$$

Suppose R fail-stop processors are required to ensure that (1)–(3) hold. Then, a computing system with $R + k$ fail-stop processors will be able to tolerate up to k fail-stop processor failures and meet its response-time goals. The obvious reconfiguration rule must be used.

Note that if stable storage that can be shared by the fail-stop processors is available, then t_{move} can be made 0. Also, by precomputing various configurations, t_{recon} can be made negligible. This, however, requires a sufficient amount of stable storage to store all possible configurations. Lastly, T_R can be made 0 by using only restartable actions; however, this uniformly degrades execution speed, even if no failures occur.

5. FAULT-TOLERANT PROCESS-CONTROL SOFTWARE

We now turn to a more substantial illustration of the application of our methodology: development of a fault-tolerant process control program. First, a correct program for a fault-free computing system is developed. The program is then extended to run correctly on a system of fail-stop processors. While a fair amount of detail is presented, these details are necessary to derive and establish the correctness of the program.

Given are *sensors* to determine the state of the environment and *actuators* to exert control over the environment. Correct operation of a process-control system requires that

PC: The values written to the actuators are related to the values read from the sensors according to a given application-specific function.

It is likely that correct operation also involves a liveness property, like “sensors are read and actuators are updated often enough.” We will make no attempt to argue that our program satisfies such real-time response constraints, although informal arguments like those developed in Section 4 could be used if timing data were available.

5.1 Assuming No Failures

Our process-control system will be structured as a collection of cyclic processes that execute concurrently. Each process p_i is responsible for controlling some set of actuators act_i . To do so, it reads from some sensors and maintains $state_i$ —a vector of *state variables* that reflects the sensor values p_i has read and the actions it has taken. Interprocess communication is accomplished by the disciplined use of shared variables; a process can read and write its state variables, but can only read state variables maintained by other processes. For the moment, we will ignore the problems that arise from concurrent access to state variables.

Each process will consist of a single loop. During execution of its *loop body*, process p_i (1) reads from some sensors, (2) computes new values for the actuators it controls and state variables it maintains, (3) writes the relevant values to act_i , and (4) updates $state_i$. Presumably, we are given application-dependent routines that can be used to compute the values to be written to the actuators and the values to be stored in the state variables.

Without loss of generality, assume that each state variable and sensor is read at most once in any execution of those routines.⁵ Let $state_j[i, t]$ denote the value of $state_j$ read by p_i during the t th execution of its loop body, $sensors[i, t]$ denote the values read by p_i from sensors during the t th execution of its loop body, and $act_i[t]$ denote the values written to act_i by p_i during the t th execution of the loop body.

Behavior satisfying PC is characterized by the following, for each process p_1, p_2, \dots, p_n .

First, the values in $state_i$ must correctly encode past actions performed by p_i . That encoding will be denoted here by the function E . Therefore, at the beginning

⁵ Code that satisfies this restriction can be written by using local variables to store state variables and sensor values: each state variable and sensor value is stored in a local variable when it is first read; subsequent references are then made to the local variable.

of the $(t + 1)$ st execution of the loop body at p_i :⁶

$$Istate(i, t): t = 0 \text{ cor } state_i = E(sensors[i, t], state_1[i, t], \dots, state_n[i, t]).$$

Secondly, the values written to actuators by p_i must be computed according to the application-specific function, here called A , based on the sensor values read and the past actions of processes. Therefore, after p_i updates act_i for the t th time,

$$Iact(i, t): t = 0 \text{ cor } act_i[t] = A(E(sensors[i, t], state_1[i, t], \dots, state_n[i, t])).$$

must be true.

Let T_i be an auxiliary variable defined so that at any time $T_i - 1$ executions of the loop body have completed. Thus, T_i is initialized to 1 and (implicitly and automatically) incremented immediately after the loop body is executed. Then, the correctness criterion PC is satisfied if

$$I(i): Istate(i, T_i - 1) \wedge Iact(i, T_i - 1)$$

is true at the beginning of each execution of the loop body, for each process p_i .

In order to construct the loop, variable *newstate* is introduced. This is necessary so that values used to update $state_i$ and the actuators are consistent with each other. Thus

$$Vnewstate(i, t): newstate = E(sensors[i, t], state_1[i, t], \dots, state_n[i, t]).$$

The loop at process p_i , which has as $I(i)$ as its loop invariant is

```

pi: process
  do true → {I(i)}
    calc: newstate := E(sensors, statei, ..., staten);
      {Vnewstate(i, Ti) ∧ Istate(i, Ti - 1) ∧ Iact(i, Ti - 1)}
    up_act: acti := A(newstate);
      {Vnewstate(i, Ti) ∧ Istate(i, Ti - 1) ∧ Iact(i, Ti)}
    up_st: statei := newstate
      {Vnewstate(i, Ti) ∧ Istate(i, Ti) ∧ Iact(i, Ti)}
  od
end

```

However, because processes execute asynchronously, access to state variables must be synchronized. Otherwise, a process might read state variables while they are in the midst of being updated, which could cause the process to perform the wrong actions. To avoid this problem, the state variables maintained by each process p_i are assumed to be characterized by CC_i , called the *consistency constraint* for $state_i$. CC_i is kept true of $state_i$ except while p_i is updating those variables, that is, performing *up_st* above. We assume that the code to compute the application dependent functions A and E works correctly as long as values that satisfy the consistency constraints are read. To ensure that only values satisfying the consistency constraints are read, read/write locks [5] can be used to implement reader-writer exclusion on the state variables maintained by each process. A process trying to read variables in $state_i$ must first acquire a read lock for $state_i$. Such a lock will not be granted if a write lock is already held for those state variables; hence that process will be delayed if $state_i$ is being updated. A

⁶ We use the notation " $A \text{ cor } B$ " to mean "if A then true else B ".

process about to update $state_i$ will be delayed if other processes are reading those values. Such lock operations are not explicitly included in our programs to simplify the exposition; they are part of the routine to compute E in $calc$ and up_st , the routine to update the state variables.

Similarly, we assume that the code to compute A and E requires that the sensor values used be consistent. The natural laws that govern our physical world ensure that at any time t the values of the sensors are consistent. Thus, if a process reads all the sensors simultaneously, consistent values would be obtained. Such a simultaneous read operation is not implementable, however. We therefore assume that sensors change values slowly enough and that processes execute quickly enough so that a consistent set of values is obtained by reading each of the sensors in sequence at normal execution speed.

5.2 Allowing Failures

We shall deal with failures by attempting to mask their effects. Thus we endeavor to preserve

PC': At no time do state variables or actuators have values they could not have had if the failure had not occurred.

Recall that $I(i)$ characterizes values of the state variables and actuators that satisfy PC. Consequently, if it is possible to modify the loop body so that $I(i)$ is true of every state that could be visible after a failure, then PC' will be satisfied, as well. Our task, therefore, is to modify the loop body so that it constitutes a restartable action.

$I(i)$ is true except from when the execution of statement up_act begins to when statement up_st completes. Thus, we must either mask intermediate states during execution of up_st and up_act , or devise a way to execute up_st and up_act together as an atomic action. This latter option is precluded by most hardware. Thus, to implement the former, we construct a single fault-tolerant action that updates the actuators and state variables on the basis of $newstate$:

```
{Vnewstate(i, Ti)}
upall
{Vnewstate(i, Ti) ∧ Istate(i, Ti) ∧ Iact(i, Ti)}
```

As long as $newstate$ is saved in stable storage, the following replete proof outline satisfies F1-F6 and accomplishes the desired transformation.

```
upall: action, recovery
      {Vnewstate(i, Ti)}
      up_act: acti := A(newstate)
      {Vnewstate(i, Ti) ∧ Iact(i, Ti)}
      up_st: statei := newstate;
      {Vnewstate(i, Ti) ∧ Istate(i, Ti) ∧ Iact(i, Ti)}
      end
```

A replete proof outline for the code executed at p_i is

```
pi: process
      action, recovery
      do true → {I(i)}
          calc: newstate := E(sensors, state1, . . . , staten);
          {Vnewstate(i, Ti) ∧ Istate(i, Ti - 1) ∧ Iact(i, Ti - 1)}
```

```

upall: action, recovery
  up__act: acti := A(newstate);
  up__st: statei := newstate;
end
od
end

```

Notice that following a failure, a process might attempt to acquire a given read/write lock that had already been granted to it. For example, if a failure occurred while *up__st* was being executed, the recovery protocol would attempt to acquire the write lock on *state_i*, which might already be owned by *p_i*. Clearly, repeated requests by a given process for the same lock, without intervening release operations, should not delay the invoker. Implementation of read/write locks with this property (binary semaphores do not suffice) is possible and is described in [20].

6. DISCUSSION

6.1 Related Work

Few general techniques have been developed to aid in the design of programs that must cope with operational failures in hardware or support software. One paradigm, based on the use of state machines, was pioneered by Lamport [12, 22]. A program is viewed as a state machine that receives input, generates actions (output), and has an internal state. A reliable system is constructed by replicating these state machines and running them in parallel. By using a solution to the Byzantine Generals Problem, each machine is guaranteed to receive the same input, despite failures. A comparison of the state machine approach with the use of fail-stop processors and fault-tolerant actions appears in [23]. A second general paradigm, which appears to be promising, is based on the use of nested atomic transactions [14].

A variety of protocols for specialized problems have also been developed. Included are protocols for recovery in database systems [5], implementation of highly reliable file systems [15], and the use of checkpoint/restart facilities in operating systems [3].

Despite the apparent similarity between the recovery block construct developed at the University of Newcastle-upon-Tyne [19] and our fault-tolerant actions, the two constructs are intended for very different purposes. A *recovery block* consists of a *primary block*, an *acceptance test*, and one or more *alternate blocks*. Upon entry to a recovery block, the primary block is executed. After its completion, the acceptance test is executed to determine if the primary block has performed acceptably. If the test is passed, the recovery block terminates. Otherwise, an alternate block—generally a different implementation of the same algorithm—is attempted, and the acceptance test is repeated. Execution of each alternate block is attempted in sequence until one produces a state in which the acceptance test succeeds. Execution of an alternate block is always begun in the recovery block's initial state.

Recovery blocks are used to mask design errors; fault-tolerant actions are used in constructing programs that must cope with *operational failures* in the underlying hardware and software. The use of recovery blocks to cope with operational

failures under such circumstances can only lead to difficulties. For example, a recovery block has only a finite number of alternate blocks associated with it, and therefore a large number of failures in the underlying system can cause the available alternatives to be exhausted. Secondly, the recovery block model does not admit the possibility of using stable storage for program variables.

6.2 Whence Fail-Stop Processors

The definition of the fail-stop processor as our underlying computational model followed from our use of a partial correctness programming logic. In a fail-stop processor all failures are detected and no incorrect state transformations result from failures. Thus, if execution of a statement terminates, by definition the transformation specified by that statement has occurred—the effect of execution is consistent with the programming logic. On the other hand, failure, by definition, prevents statements from terminating. Thus, the partial correctness (as opposed to total correctness) nature of the programming logic subsumes the consequences of failures.

6.3 Application of the Methodology

We have successfully employed the methodology described in this paper both to verify existing fault-tolerant protocols and to devise new ones. In [20], the two-phase commit protocol, as described in [5], is verified. The process-control example described in Section 5 of this paper was developed as part of a project to apply this methodology in the design of a distributed computing system for navigation in an airplane. The details of that work are discussed in [24].

It is natural to ask whether F1–F6, the components of our proof rule for fault-tolerant actions, are too restrictive. In that case there would exist fault-tolerant actions that would behave correctly, but for which no proof would be possible. While we have not proved the relative completeness of our new rule, the success we have had with its application and the way in which it was derived suggest that F1–F6 are not too restrictive to allow proof of any “correct” fault-tolerant action.

7. CONCLUSIONS

We have described a methodology for constructing fault-tolerant systems. It is based on the notion of a fail-stop processor—a processor with simple and well-defined failure-mode operating characteristics. Fail-stop processors are very appealing abstract machines to program and can be approximated by real hardware.

We have shown how axiomatic program verification techniques can be extended for proving the correctness of programs written for fail-stop processors. This allows a programmer to argue convincingly about the correctness of a program *ex post facto*. What is more important is that it allows a programmer to develop a fault-tolerant program and its proof hand-in-hand, with the latter leading the way, as advocated in [4] and [6]. Computing the weakest precondition of a recovery protocol is a simple and mechanical way of determining what program variables must be stored in stable storage; constructing a replete proof outline similarly defines what intermediate states could be visible following a failure and thus what states can be seen by a recovery protocol.

ACKNOWLEDGMENTS

This work has benefited from discussions with G. Andrews, R. W. Conway, R. Constable, E. W. Dijkstra, L. Lamport, G. Levin, R. Reitman, D. Wall, and S. Worona. The process-control application was first suggested by J. Kemp, W. Comfort, and M. Kushner of IBM (FSD/Owego). D. Gries made very helpful comments on an earlier draft of this paper. We would also like to thank the referees for a very careful reading of this paper.

REFERENCES

1. AVIZIENIS, A. Fault-tolerant systems. *IEEE Trans. Comput. C-25*, 12 (Dec. 1976), 1304-1312.
2. BARLOW, R. W., AND PROSCHAN, F. *Mathematical Theory of Reliability*. Wiley, New York, 1965.
3. DENNING, P. Fault-tolerant operating systems. *Comput. Surv.* 8, 4 (Dec. 1976), 359-389.
4. DIJKSTRA, E. W. *A Discipline of Programming*. Prentice-Hall, Englewood Cliffs, N. J., 1976.
5. GRAY, J. Notes on data base operating systems. *Operating Systems: An Advanced Course, Lecture Notes in Computer Science*, vol. 60. Springer-Verlag, New York, 1978, pp. 393-481.
6. GRIES, D. *The Science of Programming*. Springer-Verlag, New York, 1981.
7. HARTER, P., AND BERNSTEIN, A. Proving real time properties of programs with temporal logic. In *Proc. SOS-8*, Asilomar, California (Dec. 1981), 1-11.
8. HOARE, C. A. R. An axiomatic basis for computer programming. *Commun. ACM* 12, 10 (Oct. 1969), 576-580.
9. HOARE, C. A. R., AND WIRTH, N. An axiomatic definition of the programming language PASCAL. *Acta Inf.* 2 (1973), 335-355.
10. HOPKINS, A. L., SMITH, T. B., AND LALA, J. H. FTMP—A highly reliable fault-tolerant multiprocessor for aircraft. *Proc. IEEE* 66, 10 (Oct. 1978), 1221-1239.
11. LAMPORT, L. Time, clocks and the ordering of events in a distributed system. *Commun. ACM* 21, 7 (July 1978), 558-565.
12. LAMPORT, L. Using time instead of timeout for fault-tolerant distributed systems. Tech. Rep. 59, SRI Int., June 1981.
13. LAMPORT, L., SHOSTAK, R., AND PEASE, M. The Byzantine Generals Problem. *ACM Trans. Program. Lang. Syst.* 3 (July 1982) 382-401.
14. LAMPSON, B. Atomic transactions. *Distributed Systems—Architecture and Implementation, Lecture Notes in Computer Science*, vol. 105, Springer-Verlag, New York, 1981, pp. 246-265.
15. LAMPSON, B., AND STURGIS, H. Crash recovery in a distributed data storage system. To be published.
16. OWICKI, S., AND LAMPORT, L. Proving liveness properties of concurrent programs. *ACM Trans. Program. Lang. Syst.* 4, 3 (July 1982), 455-495.
17. PEASE, M., SHOSTAK, R., AND LAMPORT, L. Reaching agreements in the presence of faults. *J. ACM* 27, 2 (April 1979) 228-234.
18. PNUELI, A. The temporal semantics of concurrent programs. *Semantics of Concurrent Computation, Lecture Notes in Computer Science*, vol. 70, Springer-Verlag, New York, 1979, pp. 1-20.
19. RANDELL, B., LEE, P. A., AND TRELEAVEN, P. C. Reliability issues in computing system design. *Comput. Surv.* 10, 2 (June 1978), 123-165.
20. SCHLICHTING, R. D. *Axiomatic Verification to Enhance Software Reliability*. Ph.D. thesis, Dept. of Comput. Sci., Cornell Univ., Jan. 1982.
21. SCHLICHTING, R. D., AND SCHNEIDER, F. B. Understanding and using asynchronous message passing. In *Proc. ACM SIGACT-SIGOPS Symp. Principles of Distributed Computing* (Ottawa, Canada, Aug. 1982), ACM, New York, pp. 141-147.
22. SCHNEIDER, F. B. Synchronization in distributed programs. *ACM Trans. Program. Lang. Syst.* 4, 2 (Apr. 1982), 125-148.
23. SCHNEIDER, F. B. Fail-stop processors. *Digest of Papers from Spring CompCon '83* (San Francisco, Calif., Mar., 1983), IEEE Computer Society, New York.
24. SCHNEIDER, F. B., AND SCHLICHTING, R. D. Towards fault-tolerant process control software. In

- Proc. Eleventh Ann. Int. Symp. Fault-Tolerant Computing* (Portland, Maine, June 1981), IEEE Computer Society, New York, pp. 48-55.
25. SIEWIOREK, D., AND SWARZ, R. S. *The Theory and Practice of Reliable System Design*. Digital Press, Bedford, Mass., 1982.
 26. WENSLEY, J., WENSKY, J. H., LAMPORT, L., GOLDBERG, J., GREEN, M., LEVITT, K. N., MELLAR-SMITH, P. M., SHOSTAK, R. E., AND WEINSTOCK, C. B. SIFT: Design and analysis of a fault-tolerant computer for aircraft control. *Proc. IEEE* 66, 10 (Oct. 1978) 1240-1255.

Received November 1981; revised July 1982; Accepted February 1983