

Contents

1	The Evolution of the Recovery Block Concept	1
1.1	INTRODUCTION	1
1.2	SYSTEM STRUCTURING	2
1.3	RECOVERY BLOCKS	4
1.4	EARLY IMPLEMENTATIONS AND EXPERIMENTS	7
1.5	EXTENSIONS AND APPLICATIONS OF BASIC RECOVERY BLOCKS	9
1.6	RECOVERY IN CONCURRENT SYSTEMS	11
1.7	LINGUISTIC SUPPORT FOR SOFTWARE FAULT TOLERANCE	15
1.8	CONCLUSIONS	17

1

The Evolution of the Recovery Block Concept

BRIAN RANDELL and JIE XU

University of Newcastle upon Tyne, England

ABSTRACT

This chapter reviews the development of the recovery block approach to software fault tolerance and subsequent work based on this approach. It starts with an account of the development and implementations of the basic recovery block scheme in the early 1970s at Newcastle, and then goes on to describe work at Newcastle and elsewhere on extensions to the basic scheme, recovery in concurrent systems, and linguistic support for recovery blocks based on the use of object-oriented programming concepts.

1.1 INTRODUCTION

A research project to investigate system reliability was initiated by the first author at the University of Newcastle upon Tyne in 1971. This was at a time when the problems of software reliability had come to the fore, for example through the discussions at the 1968 and 1969 NATO Software Engineering Conferences, concerning what at the time was termed the “software crisis”. Such discussions were one of the spurs to research efforts, in a number of places, aimed at finding means of producing error-free programs. However, at Newcastle the opposite (or more accurately the complementary) problem, namely that of what to do in situations where, perhaps despite the use of the best available means of achieving error-free code, the possibility of residual design faults could not be denied, was taken as an interesting and worthwhile goal.

A preliminary phase of the project involved a study of a representative set of large software systems, including a major banking system, and an airline reservations system. This provided interesting statistical data confirming that residual software faults were one of the most important causes of system failures and down-time. It was found that in all these systems, a sizable

proportion of their code and complexity was related to provisions for (mostly hardware) fault tolerance, such as data consistency checking, and checkpointing schemes. However, these provisions, though typically rather ad hoc, were often quite effective, and indeed managed to cope with some of the software errors that were encountered in practice during system operation, even though the fault tolerance provisions had not been specifically designed to do this.

We were well aware that if we were to develop techniques aimed explicitly at tolerating software faults we would have to allow for the fact that the principal cause of residual software design faults is complexity. Therefore the use of appropriate structuring techniques would be crucial — otherwise the additional software that would be needed might well increase the system's complexity to the point of being counter-productive. Aided by what we had found in our examination of the checkpoint and restart facilities then being employed, we came to realize that although a variety of even quite disparate error detection mechanisms could usefully be employed together in a system, it was critical to have a simple, coherent and general strategy for error recovery. Moreover it was evident that such a strategy ought to be capable of coping with multiple errors, including ones that were detected during the error recovery process itself.

The first structuring technique that we developed was in fact the basic “recovery block” scheme. In what follows we use the structuring concepts that we later developed in our description of this basic scheme, and of some of the ensuing research on recovery blocks carried out at Newcastle and elsewhere, before discussing some of the latest ideas that we have been investigating on the structuring of fault-tolerant software.

1.2 SYSTEM STRUCTURING

Our interest in the problems of structuring systems so as to control their complexity, and in particular that of their fault tolerance provisions, led us to a style of system design which is based on what we term *idealized fault-tolerant components* [And81, Ran84]. Such components provide a means of system structuring which makes it easy to identify *what* parts of a system have *what* responsibilities for trying to cope with *which* sorts of fault.

We view a system as a set of components interacting under the control of a design (which is itself a component of the system) [Lee90]. Clearly, the system model is recursive in that each component can itself be considered as a system in its own right and thus may have an internal design which can identify further sub-components. Components receive requests for service and produce responses. When a component cannot satisfy a request for service, it will return an exception. An idealized fault-tolerant component should in general provide both normal and abnormal (i.e. exception) responses in the interface between interacting components, in a framework which minimizes the impact of these provisions on system complexity. Three classes of exceptional situation (i.e. in which some fault tolerance response is needed) are identified. Interface exceptions are signaled when interface checks find that an invalid service request has been made to a component. These exceptions must be treated by the part of the system which made the invalid request. Local exceptions are raised when a component detects an error that its own fault tolerance capabilities could or should deal with in the hope that the component would return to normal operations after exception handling. Lastly, a failure exception is signaled to notify the component which made the service request that, despite the

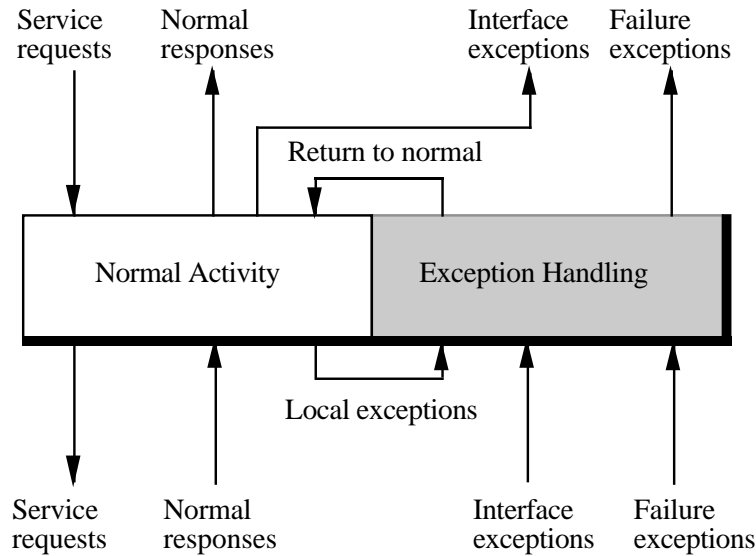


Figure 1.1 Idealized component

use of its own fault tolerance capabilities, it has been unable to provide the service requested of it (see Figure 1.1).

Our notion of an idealized component is mainly concerned with interactions of a component with its environment. It makes minimal assumptions on faults, fault masking and the fault tolerance scheme adopted, in indicating how exception-handling should be structured. Exception handling is often considered as being a limited form of software fault tolerance — for example, by detecting and recovering an error, and either ignoring the operation which generated it or by providing a pre-defined and heavily degraded response to that operation. However, such software cannot be regarded as truly fault-tolerant since some perceived departure from specification is likely to occur, although the exception handling approach can result in software which is robust in the sense that catastrophic failure can often be avoided.

In order also to achieve effective design fault tolerance, capable of completely masking the effects of many residual software errors, it is necessary to incorporate deliberate redundancy, i.e. to make use of design diversity, in such systems. The structuring scheme that we have developed [Ran93] both for describing and comparing the various existing software fault tolerance schemes, and indeed for guiding their actual implementation, is illustrated in Figure 1.2.

This shows an idealized component which consists of several sub-components, namely an adjudicator and a set of software variants (modules of differing design aimed at a common specification). The design of the component, i.e. the algorithm which is responsible for defining the interactions between the sub-components, and establishing connections between the component and the system environment, is embodied in the controller. This invokes one or more of the variants, waits as necessary for such variants to complete their execution and invokes the adjudicator to check on the results produced by the variants. As illustrated in Figure 1.2, each of these sub-components (even the adjudicator), as well as the component (controller) itself, can in principle contain its own provisions for exception handling, and indeed for full software fault tolerance, so the structuring scheme is fully recursive.

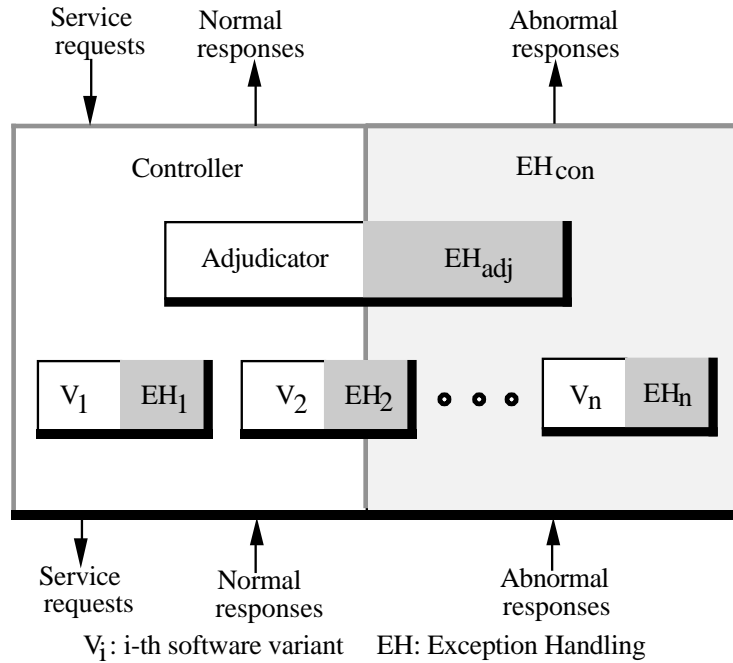


Figure 1.2 Idealized component with design redundancy

Obviously, the notion of structuring systems needs to be used in such a way as to achieve an appropriate structuring of the complex asynchronous activities to which the system can give rise, in particular those related to fault tolerance. In common with other groups, we make use of so-called *atomic actions* for this purpose. The activity of a group of components constitutes an atomic action if no information flows between that group and the rest of the system for the duration of the activity [Lee90]. Atomic actions may be planned when the system is designed, or (less commonly) may be dynamically identified by exploratory techniques after the detection of an error. Planned atomic actions must be maintained by imposing constraints on communication within the system. Error recovery can be linked to the notion of an atomic action, which is said to form a restorable action if all components within the action retain the ability to perform a mutually consistent state restoration. These issues are discussed further in Section 1.6.

1.3 RECOVERY BLOCKS

In this section, we discuss recovery blocks in detail, making use of the exception handling terminology introduced above. The basic recovery block relates to sequential systems. Details of extensions for use in concurrent systems are discussed in Section 1.6. The recovery block approach attempts to prevent residual software faults from impacting on the system environment, and it is aimed at providing fault-tolerant functional components which may be nested within a sequential program. The usual syntax is as follows:

```

ensure          acceptance test
by             primary alternate
else by        alternate 2
              .
              .
else by        alternate n
else error

```

Here the alternates correspond to the variants of Figure 1.2, and the acceptance test to the adjudicator, with the text above being in effect an expression of the controller. On entry to a recovery block the state of the system must be saved to permit backward error recovery, i.e. establish a checkpoint. The primary alternate is executed and then the acceptance test is evaluated to provide an adjudication on the outcome of this primary alternate. If the acceptance test is passed then the outcome is regarded as successful and the recovery block can be exited, discarding the information on the state of the system taken on entry (i.e. checkpoint). However, if the test fails or if any errors are detected by other means during the execution of the alternate, then an exception is raised and backward error recovery is invoked. This restores the state of the system to what it was on entry. After such recovery, the next alternate is executed and then the acceptance test is applied again. This sequence continues until either an acceptance test is passed or all alternates have failed the acceptance test. If all the alternates either fail the test or result in an exception (due to an internal error being detected), a failure exception will be signaled to the environment of the recovery block. Since recovery blocks can be nested, then the raising of such an exception from an inner recovery block would invoke recovery in the enclosing block. The operation of the recovery block is further illustrated in Figure 1.3.

Obviously, the linguistic structure for recovery blocks requires a suitable mechanism for providing automatic backward error recovery. Randell produced the first such "recovery cache" scheme, a description of which was included in the first paper on recovery blocks [Hor74] (although this scheme was later superseded [And76]). This paper also included a discussion of "recoverable procedures" — a rather complex mechanism that Lauer and Randell had proposed as a means of extending the recovery cache scheme to deal with programmer-defined data types. This part of the paper would undoubtedly have been much clearer if the ideas had been expressed in object-oriented terms — a point we will develop further in Section 1.7.

The overall success of the recovery block scheme rests to a great extent on the effectiveness of the error detection mechanisms used — especially (but not solely) the acceptance test. The acceptance test must be simple otherwise there will be a significant chance that it will itself contain design faults, and so fail to detect some errors, and/or falsely identify some conditions as being erroneous. Moreover, the test will introduce a run-time overhead which could be unacceptable if it is very complex. The development of simple, effective acceptance tests can thus be a difficult task, depending on the actual specification.

In fact, the acceptance test in a recovery block should be regarded as a last line of detecting errors, rather than the sole means of error detection. The expectation is that it will be buttressed by executable assertion statements within the alternates and run-time checks supported by the hardware. Generally, any such exception raised during the execution of an alternate will lead to the same recovery action as for acceptance test failure. Should the final alternate fail, for example by not passing the acceptance test, this will constitute a failure of the entire

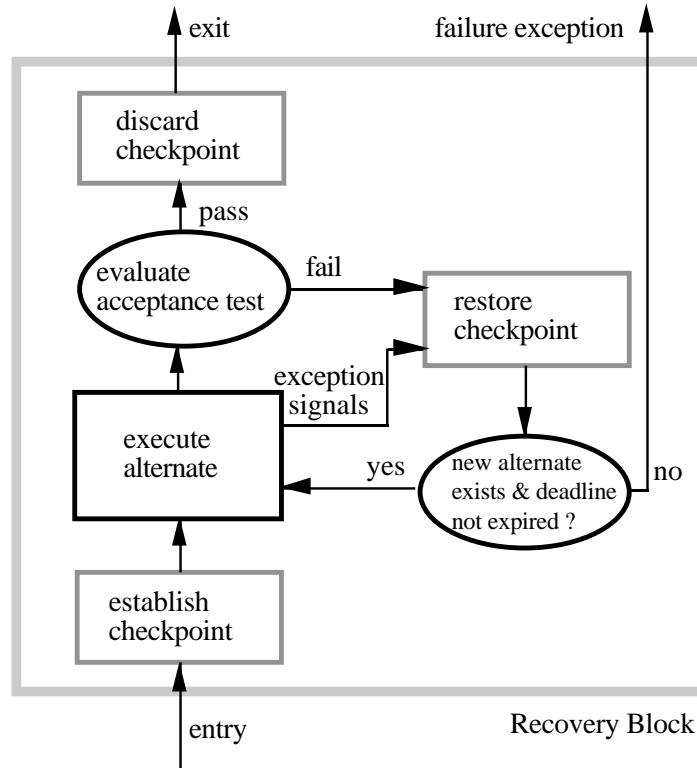


Figure 1.3 Operation of the recovery block

module containing the recovery block, and will invoke recovery at the level of the surrounding recovery block, should there be one.

In other words, each alternate should itself be an ideal fault-tolerant component. An exception raised by run-time assertion statements within the alternate or by hardware error-detection mechanisms may be treated by the alternate's own fault tolerance capabilities. A failure exception is raised to notify the system (i.e. the control component in our model) if, despite the use of its own fault tolerance capabilities, the alternate has been unable to provide the service requested of it. The control component may invoke then another alternate.

In general, as described in [Mel77], forward error recovery can be further incorporated in recovery blocks to complement the underlying backward error recovery. (In fact, a forward error recovery mechanism can support the implementation of backward error recovery by transforming unexpected errors into default error conditions [Cri82].) If, for example, a real-time program communicated with its (unrecoverable) environment from within a recovery block then, if recovery were invoked, the environment would not be able to recover along with the program and the system would be left in an inconsistent state. In this case, forward recovery would help return the system to a consistent state by sending the environment a message informing it to disregard previous output from the program.

In the first paper about recovery blocks [Hor74], Horning *et al.* list four possible failure conditions for an alternate: i) failure of the acceptance test, ii) failure to terminate, detected by a timeout, iii) implicit error detection (for example divide by zero), and iv) failure ex-

ception of an inner recovery block. Although the mechanism for implementing the time-out detection measure was not discussed by the authors, the original definition of recovery blocks does cover this issue. Several implementations of watchdog timers for recovery blocks have been described [Hec76, Kim89]. Timeout can be provided as a syntactic form in the recovery block structure [Gre85]. As with the `else error` part, the `timeout` part allows the programming of a "last ditch" algorithm for the block to achieve its goal, and is really a form of forward recovery since its effects will not be undone (at least at this level).

Although each of the alternates within a recovery block endeavors to satisfy the same acceptance test there is no requirement that they all must produce the same results [Lee78]. The only constraint is that the results must be acceptable — as determined by the test. Thus, while the primary alternate should attempt to produce the desired outcome, the further alternate may only attempt to provide a degraded service. This is particularly useful in real-time systems, since there may be insufficient time available for fully-functional alternates to be executed when a fault is encountered. An extreme corresponds to a recovery block which contains a primary module and a null alternate [And83, And85]. Under these conditions, the role of the recovery block is simply to detect and recover from errors by ignoring the operation where the fault manifested itself. This is somewhat similar to forward error recovery because the manifestation of a fault will result in a loss of service to the environment. But the important difference is that forward recovery can only remove predictable errors from the system state, whereas such backward recovery can still cope with the unpredictable errors caused by residual design faults. (The only constraint is that the errors do not impact the actual recovery block mechanism.)

Most of the time, only the primary alternate of the recovery block is executed. (This keeps the run-time overhead of the recovery block to a minimum and makes good use of the system and hardware resources.) However, this could cause a problem: the alternates must not retain data locally between calls, otherwise these modules could become inconsistent with each other since not all of them are executed each time when the recovery block is invoked. The problem becomes more obvious while one attempts to design an alternate as an object. There is no guarantee that the state of the object is correctly modified unless the object is invoked each time. Distributed (parallel) execution of recovery blocks [Kim84] could solve this issue. An alternative solution is to design the alternate modules as memoryless functional components rather than as objects.

Unlike tolerance to hardware malfunctions, software fault tolerance cannot be made totally transparent to the application programmer although some operations related to its provision, such as saving and restoring the state of the system, can be made automatic and transparent. The programmer who wishes to use software fault tolerance schemes must provide software variants and adjudicators. Therefore, a set of special linguistic features or conventions is necessary for incorporating software redundancy in programs. The key point here is to attempt to keep the syntactic extension simple, natural and minimal. This will be further discussed in the Section 1.7.

1.4 EARLY IMPLEMENTATIONS AND EXPERIMENTS

The first implementation of recovery blocks involved defining and simulating a simple stack-oriented instruction set, incorporating a recovery cache [And76]. Simple test programs embodying recovery blocks could be run on this machine simulator, and subjected to deliberate

faults. Test programs were run on one computer — a separate computer was used to provide data to, and to accept and check the output from, these programs. This second computer also provided facilities by means of which experimenters could make arbitrary changes to any locations in the simulated memory. Visitors to the project were typically challenged to use these facilities to try and cause a demonstration recovery block program to fail — their inability to do so was a persuasive argument for the potential of the recovery block scheme!

Another experimental system is described in [Shr78a, Shr78b] in which recovery blocks were incorporated in the language Pascal. The modification was made to the kernel and interpreter of Brinch Hansen's Pascal system to support the syntax of recovery blocks and the associated recovery caches needed for state restoration. Based on this extension and on a few experimental programs, some performance measurements for recovery blocks were reported, which generally support the belief that recovery blocks do not impose any serious runtime and recovery data space overheads. For the sample programs, the run-time overhead ranged between 1 to about 11% of T1 (execution time of a program without any recovery facilities) when no errors are detected. When a primary fails, the time taken to restore system state was up to about 30% of T1. This experiment also showed that recovery caches made a substantial saving in space, compared with complete checkpointing.

The recovery cache mechanism should ideally form an integral part of a given computer; this not being possible for the existing hardware. The next major work at Newcastle on the implementation of the basic recovery block scheme involved the design and building of a hardware recovery cache for the PDP-11 family of machines [Lee80]. This device was inserted into the bus between the CPU and memory modules without requiring hardware alterations. It intercepted writes to memory, and automatically determined whether the contents of the memory location that was about to be over-written needed to be saved beforehand. In order to minimize the overheads imposed on the host, special hardware was designed to enable concurrent operation of the recovery cache and the host system.

The controversial nature of software fault tolerance spurred extensive efforts aimed at providing evidence of the scheme's potential cost-effectiveness in real systems. (The developers of *N*-version programming [Avi77] were similarly motivated to undertake extensive experimental evaluations, as discussed in Section 2.4 in Chapter 2.) During 1981-84 therefore, a major project directed by Tom Anderson applied an extension of recovery blocks in the implementation of a Naval Command and Control system composed of about 8000 lines of CORAL programming, and made use of the above-mentioned hardware cache [And85]. The practical development work of the project included the design and implementation of a virtual machine which supported recovery blocks, together with extensions to the CORAL programming language to allow software fault-tolerance applications to be written in this high-level language. To maintain realism the system was constructed by experienced programmers in strict accordance with the official rules for defense-related software projects. Analysis of experimental runs of this system showed that a failure coverage of over 70% was achieved. The supplementary cost of developing the fault-tolerant software was put at 60% of the implementation cost. The system overheads were measured at 33% extra code memory, 35% extra data memory and 40% additional run time. These led to the conclusion that "by means of software fault tolerance a significant and worthwhile improvement in reliability can be achieved at acceptable cost" [And85].

Research at the Royal Military College of Science subsequently extended this experiment to the design of a demonstrator modeled on functions provided at the London Air Traffic

Control Center, and the results have reinforced confidence in the cost-effectiveness and the general applicability of the recovery block approach [Mou87].

1.5 EXTENSIONS AND APPLICATIONS OF BASIC RECOVERY BLOCKS

Many applications and varieties of recovery blocks have been explored and developed by various researchers. Some of typical experiments and extensions are considered below.

1.5.1 Distributed Execution of Recovery Blocks

H. Hecht was the first to propose the application of recovery blocks to flight control systems [Hec76, Hec86]. His work included an implementation of a watchdog timer that monitors availability of output within a specified time interval and his model also incorporates a rudimentary system to be used when all alternates of the recovery block scheme are exhausted. Since then, further researches and experiments have been conducted by Hecht and his colleagues. For example, M. Hecht *et al.* [Hec89] described a distributed fault-tolerant architecture, called the extended distributed recovery block, for nuclear reactor control and safety functions. Their architecture relies on commercially available components and thus allows for continuous and inexpensive system enhancement. The fault injection experiments during the development process demonstrate that the system could tolerate most single faults and dual faults.

K. H. Kim and his colleagues in the DREAM Laboratory have extensively explored the concept of distributed execution of recovery blocks, a combination of both distributed processing and recovery blocks, as an approach for uniform treatment of hardware and software faults [Kim84, Kim89, Kim88b, Wel83]. The details are given in Chapter 8. A useful feature of their approach is the relatively low run-time overhead it requires so that it is suitable for incorporation into real-time systems. The basic structure of the distributed recovery block is straightforward: the entire recovery block, two alternates with an acceptance test, is fully replicated on the primary and backup hardware nodes. However, the roles of the two alternate modules are not the same in the two nodes. The primary node uses the first alternate as the primary initially, whereas the backup node uses the second alternate as the initial primary. Outside of the distributed recovery block, forward recovery can be achieved in effect; but the node affected by a fault must invoke backward recovery by executing an alternate for data consistency with the other nodes. To test the execution efficiency of the approach, two experimental implementations and measurements have been conducted on distributed computer networks. The results indicate the feasibility of attaining fault tolerance in a broad range of real-time applications by means of the distributed recovery blocks.

1.5.2 Consensus Recovery Blocks

The consensus recovery block (CRB) [Sco85] is an attempt to combine the techniques used in the recovery block and N -version programming [Avi77]. It is claimed that the CRB technique reduces the importance of the acceptance test used in the recovery block and is able to handle the case where NVP would not be appropriate since there are multiple correct outputs. The CRB requires design and implementation of N variants of the algorithm which are ranked (as in the recovery block) in the order of service and reliance. On invocation, all variants

are executed and their results submitted to an adjudicator, i.e. a voter (as used in N -version programming). The CRB compares pairs of results for compatibility. If two results are the same then the result is used as the output. If no pair can be found then the results of the variant with the highest ranking are submitted to an acceptance test. If this fails then the next variant is selected. This continues until all variants are exhausted or one passes the acceptance test.

[Sco87] developed reliability models for the recovery block, N -version programming and the CRB. In comparison, the CRB is shown to be superior to the other two. However, the CRB is largely based on the assumption that there are no common faults between the variants. (This of course is not the case, as was shown by such experiments as [Kni85, Sco84].) In particular, if a matching pair is found, there is no indication that the result is submitted to the acceptance test, so a correlated failure in two variants could result in an erroneous output and would cause a catastrophic failure.

1.5.3 Retry Blocks with Data Diversity

A retry block developed by Ammann and Knight [Amm87, Amm88] is a modification of the recovery block scheme that uses data diversity instead of design diversity. Data diversity is a strategy that does not change the algorithm of the system (just retry), but does change the data that the algorithm processes. It is assumed that there are certain data which will cause the algorithm to fail, and that if the data were re-expressed in a different, equivalent (or near equivalent) form the algorithm would function correctly. A retry block executes the single algorithm normally and evaluates the acceptance test. If the test passes, the retry block is complete. If the test fails, the algorithm executes again after the data has been re-expressed. The system repeats this process until it violates a deadline or produces a satisfactory output. The crucial elements in the retry scheme are the acceptance test and the data re-expression routine.

A description of some experiments with the retry block is presented by the authors. Coordinates to a radar system were altered to lie on the circumference of a small circle centered on the point, taking advantage of the fact that this application's data had limited precision. The radius of the circle and the re-expression algorithm were both changed to generate an indication of their influence. Although the overall performance of the retry block varied greatly, a large reduction in failure probability for some of the faults is observed in the study. Compared with design diversity, data diversity is relatively easy and inexpensive to implement. Although additional costs are incurred in the algorithm for data re-expression, data diversity requires only a single implementation of a specification. Of course, the retry scheme is not generally applicable and the re-expression algorithm must be tailored to the individual problem at hand and should itself be simple enough to eliminate the chance of design faults.

1.5.4 Self-Configuring Optimal Programming

SCOP (Self-configuring optimal programming) [Bon93, Xu93] is another attempt to combine some techniques used in RB and NVP in order to enhance efficiency of software fault tolerance and to eliminate some inflexibilities and rigidities. This scheme organizes the execution of software variants in phases, dynamically configuring a currently active variant set, so as to produce acceptable results with the relatively small effort and to make the efficient use of available resources. The control can be parameterized with respect to the level of fault toler-

ance, the amount of available resources and the desired response time. Since highly dynamic behavior can cause complexity of control and monitoring, a methodology for devising various instances of SCOP is developed by simplifying the on-line process at the price of the complex off-line design.

The gain of efficiency would be limited when the supporting system is intended for a specific application – the hardware resources saved by the SCOP scheme would be merely left idle. It is perhaps more appropriate if the application environments are complex and highly variable, such as a large distributed computing system that supports multiple competing applications.

1.5.5 Other Applications

Sullivan and Masson developed an algorithm-oriented scheme, based on the use of what they term Certification Trails [Sul90, Sul91]. The central idea of their method is to execute an algorithm so that it leaves behind a trail of data (certification trail) and, by using this data, to execute another algorithm for solving the same problem more quickly. The outputs of the two executions are compared and considered correct only if they agree. An issue with the data trail is that the first algorithm may propagate an error to the second algorithm, and this could result in an erroneous output. Nevertheless, the scheme is an interesting alternative to the recovery block scheme, despite being perhaps of somewhat limited applicability.

Delta-4 was a collaborative project carried out within the framework of the European Strategic Program for Research in Information Technology (ESPRIT) [Pow91]. Its aim was the definition and design of an open, dependable, distributed computer system architecture. The Delta-4 approach deals mainly with hardware fault tolerance, but also addresses the issue of design faults. [Bar93] describes the integration of software fault tolerance mechanisms into the existing Delta-4 architecture. The authors claimed that the incorporation of recovery blocks and dialogues (structures for supporting inter-process recovery) into the Delta-4 framework is obtained without significant overheads.

1.6 RECOVERY IN CONCURRENT SYSTEMS

Work at Newcastle on this topic dates from 1975, when we began to consider the problems of providing structuring for error recovery among sets of cooperating processes. (A few researches were also made into error recovery in the particular case of so-called competing processes where the processes communicate only for resource sharing [Shr78c, Shr79].) Having identified the dangers of what we came to term the *domino effect*, we came up with the notion of a *conversation* [Ran75] — something which we later realized was a special case of a nested atomic action.

1.6.1 Conversations

When a system of cooperating processes employs recovery blocks, each process will be continually establishing and discarding checkpoints, and may also need to restore to a previously established checkpoint. However, if recovery and communication operations are not performed in a coordinated fashion, then the rollback of a process can result in a cascade of rollbacks that could push all the processes back to their beginnings — the domino effect. This

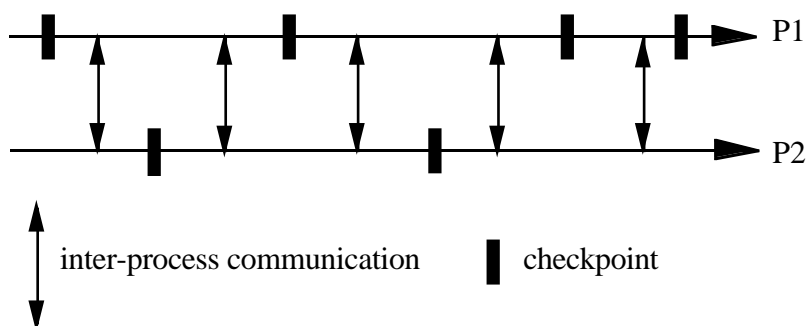


Figure 1.4 The domino effect

causes the loss of entire computation performed prior to the detection of the error. Figure 1.4 illustrates the domino effect with two communicating processes.

The conversation scheme is in our view one of the fundamental approaches to structured design of fault-tolerant concurrent programs. It provides a means of coordinating the recovery blocks of interacting processes to avoid the domino effect. Figure 1.5 shows an example where three processes communicate within a conversation and the processes P1 and P2 communicate within a nested conversation. Communication can only take place between processes that are participating in a conversation together. The operation of a conversation is: (i) on entry to a conversation a process establishes a checkpoint; (ii) if an error is detected by any process then all the participating processes must restore their checkpoints; (iii) after restoration all processes use their next alternates; and (iv) all processes leave the conversation together. The concept of conversation facilitates failure atomicity and backward recovery in cooperating process systems in a manner analogous to that of the atomic action mechanism in object-based systems. In fact, this terminological distinction between the area of communicating process systems and that of object-based systems is, we claim, of only surface importance [Shr93].

Considerable research has been undertaken into the subject of concurrent error recovery, including improvements on the conversation and different implementations of it. There are at least two classes of approaches to preventing the domino effect: the coordination-by-programmer approach and the coordination-by-machine approach. With the first approach, the application programmer is fully responsible for designing processes so that they establish checkpoints in a well coordinated manner [Ran75, Rus80, Kim82]. Many authors have added language constructs to facilitate the definition of restorable actions based on this approach [Rus79, And83, Gre85, Jal84, Jal86]. In contrast, the coordination-by-machine approach relies on an "intelligent" underlying processor system which automatically establishes appropriate checkpoints of interacting processes [Kim78, Bar83, Koo87, Kim90]. If restorable actions are unplanned, so that the recovery mechanism must search for a consistent set of checkpoints, such actions would be expensive and difficult to implement. However, such exploratory techniques have the advantage that no restrictions are placed on inter-process communication and that a general mechanism could be applied to many different systems [Mer78, Woo81]. To reduce synchronization delays introduced by controlled recovery, some researches have focused on the improvement of performance, such as the lookahead scheme and the pseudo-recovery block [Kim76, Kim88a, Ram88, Rus79, Shi84].

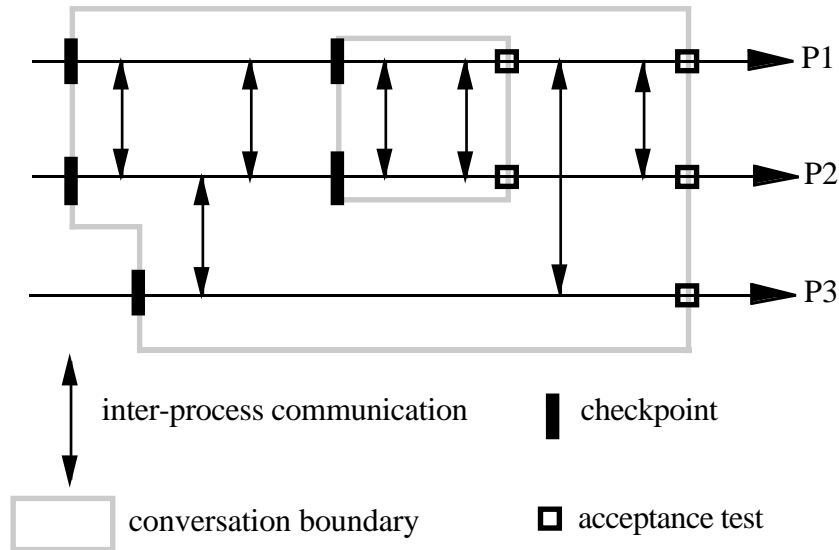


Figure 1.5 Nested conversations

1.6.2 Extensions and Implementations of Conversations

The original description of conversations provided a structuring or design concept without any suggested syntax. [Rus79] proposed a syntax called the name-linked recovery block for the concept of conversations. Kim [Kim82] presents three different syntactic forms for conversations based on the monitor structure. The different implementations deal with the distribution of the code for the recovery blocks of individual processes. The tradeoff is either to spread the conversation among the individual processes such that all of the code of each process is in one location or have all the code for the conversation in one location.

There was no provision for linked forward error recovery in the original conversation scheme. Campbell and Randell [Cam86] proposed techniques for structuring forward error recovery measures in asynchronous systems and generalized ideas of atomic actions so as to support fault-tolerant interactions between processes. A resolution scheme is used to combine multiple exceptions into a single exception if they are raised at the same time.

Issarny extended their work to concurrent object-oriented systems by defining an exception-handling mechanism for parallel object-oriented programming [Iss93a]. This mechanism was then generalized to support both forward and backward error recovery [Iss93b]. Following the proposal in [Cam86], Jalote and Campbell described a system which contains both forward and backward error recovery within a conversation structure (also known as an FT-action). Their system was based on communicating sequential processes (CSP) [Hoa78] with one extension (the `exit`) statement.

Forward error recovery in an FT-action [Jal86] is achieved through linked exception handlers where each process has its own handler for each exception. When an exception is raised by a process it is propagated to all the participating processes within the FT-action. Each process then executes its own handler for that exception. Backward recovery within an FT-action is obtained by recovery blocks. Every participating process is required to have the same number of alternates. An FT-action can combine the two schemes so that forward and backward

error recovery are used within the same structure. It can also cope with the issue of real-time applications through a simple timer.

Real-time applications may suffer from the possibility of *deserters* in a conversation — if a deadline is to be met then a process that fails to enter the conversation or to reach its acceptance test could cause all the processes in the conversation to miss that deadline [Kim82]. Russell and Tiedeman [Rus79] considered relaxing the requirement for all processes exiting together so as to enable some protection against deserter processes, but this could lead to the domino effect. Campbell, Horton and Belford [Cam79] proposed a deadline mechanism for dealing with timing faults. Anderson and Knight [And83] proposed *exchanges* as a simplification of conversations where the cyclic nature of real-time systems is exploited. (An exchange is a conversation in which all participating processes enter upon initiation and terminate upon exit. Error recovery is particularly easy as the recovery data is only that needed upon initiation, which should only be a small amount of frame dependent data.)

Gregory and Knight [Gre85] identified a set of problems associated with conversations. They argued that there ought to be two types of acceptance test — one for each process within a conversation to check its own goal and one for the whole structure of the conversation to check the global goal. In addition, within a conversation or other structures mentioned above the set of processes that attempt their primary alternate is the same as the set of processes which attempt all other alternates, i.e. they all roll back and try again with their further alternates. This is overly restrictive and affects independence of algorithm between alternates. In an effort to solve these problems, they developed two concepts — a *colloquy* that contains many *dialogs*.

A dialog is a way of enclosing a set of processes in an atomic action. It provides no retry method and no definition of the action to be taken upon failure. If any failure occurs, the dialog restores all checkpoints and fails, signaling the failure to the surrounding colloquy. A colloquy that contains a set of dialogs controls the execution of dialogs and decides on the recovery action to be taken if the dialog fails. The colloquy provides a means of constructing alternates using a potentially different set of processes, thereby permitting true diverse design. The dialog and colloquy allow time constraints to be specified and are accompanied by syntactic proposals that are extensions to the Ada language.

However, when attempting the integration of the syntax for the colloquy into Ada, the authors found several new and potentially serious difficulties which arise because of a conflict between the semantics of modern programming languages and the needs of concurrent backward recovery [Gre89].

1.6.3 Practical Difficulties and Possible Solutions

The practical problems mentioned in [Gre89] fall into the general categories of (i) program structure, (ii) shared objects, and (iii) process manipulation. All the problems have the potential to allow the state outside a dialog (or a conversation) to be contaminated by changes inside the dialog, i.e. *information smuggling*. For example, a major inconsistency exists between the preferred structure of concurrent programs (e.g. involving the use of service processes) and the structure for planned recovery. To avoid information smuggling, the planned backward recovery could cause the *capture effect* for service processes — in other words processes outside the (nested) dialogs cannot use the service processes until the completion of all the dialogs. Shared objects are another significant source of information smuggling, but no simple approaches solve the problem. Smuggling can occur with process manipulations (e.g. dynamic

process creation) also. An initial solution to the problem merely raises several other issues. Given the complexity and subtlety of these problems, Gregory and Knight concluded that "the only workable solution might be that programming language design begin with backward error recovery as its starting point." Nevertheless, some preliminary and partial solutions can be found in [Cle93, Gre87].

The actual programming of a conversation is another major difficulty associated with the concept. Constructing an application into a sequence of conversations is not a trivial task. The application programmer has to select a boundary composed of a set of checkpoints, acceptance tests and the side walls to prevent information smuggling. This boundary should be integrated well into the structure of processes. [Tyr86] suggested a way of identifying adequate boundaries of conversations based on the specification of the application using Petri Nets. [Car91] proposed an alternative solution in which the CSP notation [Hoa78] is used to describe the application and conversation boundaries are identified through a trace evaluation, but such traces would cause an explosion of states even for simple applications. In practice, however, it is possible for some special applications to decide on the conversation placement without full trace evaluation [Tyr92].

1.7 LINGUISTIC SUPPORT FOR SOFTWARE FAULT TOLERANCE

General linguistic supports for software fault tolerance are the concern of much of our latest work. If the design of software fault-tolerant systems is to become widely used on a routine basis, one of important problems that has to be solved is how to develop and provide appropriate linguistic notations and the corresponding environments, which should effectively support the development of fault-tolerant programs without greatly complicating the program's implementation, readability, and maintenance.

1.7.1 Design Notations and Environments

[Liu92] proposed a design notation for a wide class of fault-tolerant software structures, mainly offering generality and flexibility in a modular fashion. [Bon92] showed that their BSM design description language is sufficient for expressing the typical structures of software fault tolerance, such as recovery blocks and N -version programming, without requiring semantic extensions. [Anc90] described a mechanism, called the *Recovery Metaprogram* (RMP), for the incorporation of fault tolerance functions into application programs, giving programmers a single environment that lets them use the appropriate fault tolerance scheme.

The architecture proposed in [Anc90] contains three components: the application program, the RMP and the kernel. The application programmer must define the software variants and the validation test, and indicate which portions of the application program are involved in fault tolerance. The RMP implements the controllers and the supporting mechanisms for four different schemes, inserting a number of breakpoints in the program. When a breakpoint is reached, the application program is suspended and the kernel activates the RMP which takes actions to support the fault tolerance scheme chosen. The RMP is then suspended, and the application program is reactivated until the next breakpoint is reached. The implementation of the RMP approach may incur an additional cost in the form of intensive context switches and kernel calls.

However, in contrast to the languages and environments discussed above, our major work

has been greatly influenced by the now very fashionable topic of object-oriented programming. In particular, we have found it convenient to try to exploit various characteristics of C++ [Str91], a language that has been used extensively at Newcastle in connection with work on distributed systems [Shr91].

1.7.2 Implementing Software Fault Tolerance in C++

The recent extension of C++ to include generic classes and functions (“templates”), and exception handling (“catch” and “throw”) makes it possible to implement both forward and backward error recovery in C++ in the form of reusable components that separate the functionality of the application from its fault tolerance [Rub93]. More generally, such facilities show prospect of providing a convenient means of achieving high levels of reuse. This would apply both to general software components implementing various fault tolerance strategies (including generalizations and combinations of recovery blocks, and N -version programs, and encompassing the use of parallelism) and to application-specific software components [Ran93]. We also provided a set of pre-defined C++ classes to support a general object-oriented framework for software fault tolerance based directly on the abstract model represented by Figure 1.2, and described in Section 1.2 [Xu94]. However, there remain certain strategies and types of structuring that cannot be implemented entirely (or at any rate elegantly) in a language like C++ even given such mechanisms as generic functions and inheritance. Instead, the programmer who wishes to employ these strategies has to obey certain conventions. For example, the application programmer who wishes to make use of our C++ classes would have to include explicit calls in each operation of an object to facilities related to the provision of state restoration.

Adherence to such conventions can be automated, by embodying them into a somewhat enhanced version of C++ and using a pre-processor to generate conventional C++ programs automatically. Although the pre-processor approach can be quite practical it does have disadvantages. In particular the language provided to application programmers becomes non-standard since programmers have in some circumstances during program development to work in terms of the program generated by the pre-processor, rather than of the one that they had written. The alternative, that of leaving it to the programmer to adhere to the conventions, is of course a fruitful source of residual program faults. But developing a new language that provides adequate syntax and runtime support to enable the implementation of various software fault tolerance could cut the work off from the mainstream of programming language developments and thus have difficulty in achieving wide acceptance.

1.7.3 Reflection and Reflective Languages

As mentioned in Section 1.3, it has to be the responsibility of the application programmers for developing software variants, acceptance tests, and even application-specific voters. Special language features and/or programming conventions therefore cannot be avoided completely. In consideration of software reliability, the key problem would become how a set of simple (thus easy to check) programming features can be developed with powerful expressibility to enable the implementation of software fault tolerance and how the supporting mechanisms can be provided in a more natural and modular manner rather than by an ad-hoc method such as system calls. Recent developments in the object-oriented language world, under the term “*reflection*” [Mae87], show considerable promise in this regard.

A reflective system can reason about, and manipulate, a representation of its own behavior. This representation is called the system's meta-level [Agh92]. Reflection improves the effectiveness of the object-level (or base-level) computation by dynamically modifying the internal organization (actually the meta-level representation) of the system so as to provide powerful expressibility. Therefore, in a reflective programming language a set of simple, well-defined language features could be used to define much more complex, dynamically changeable constructs and functionalities. In our case, it could enable the dynamic change and extension of the semantics of those programming features that support software fault tolerance concepts, whereas the application-level (or object-level) program is kept simple and elegant [Xu94]. Although C++ itself does not provide a metalevel interface, Chiba and Masuda [Chi93] describes an extension of C++ to provide a limited form of computational reflection, called Open C++, whose usefulness in expressing software fault tolerance we are now investigating.

However, quite what reflective capabilities are needed for what forms of fault tolerance, and to what extent these capabilities can be provided in more-or-less conventional programming languages, and allied to the other structuring techniques outlined in this chapter, remain to be determined. In particular, the problems of the combined provision of significant software fault tolerance and hardware fault tolerance, and of evaluating cost-effectiveness, are likely to require much further effort. When considering support for software fault tolerance in concurrent object-oriented programming, we face a greater challenge because, on the one hand, mainstream object-oriented languages such as C++ and Eiffel [Mey92] do not at present address concurrency and, on the other hand, a large number of different models for concurrent object-oriented programming have been proposed but none has yet received widespread acceptance. There exist only a few tentative proposals for treating concurrent error recovery such as the Arche language [Ben92]. However, the reflection technique seems to be a more promising approach to the structuring of concurrent object-oriented programs [Yon89].

1.8 CONCLUSIONS

Looking back on the developments that have occurred since the recovery block concept was first introduced, it is we hope fair to claim that it has proved a very useful abstraction, and starting point for much subsequent research, elsewhere as well as at Newcastle. That at Newcastle can be characterized as mainly involving over the years:

- a gradual extension of the original very basic scheme to deal with ever more complex situations, while retaining as much as possible of the essential simplicity of structuring provided by the basic scheme, and more recently (and perhaps rather belatedly)
- the investigation of appropriate linguistic support for recovery blocks and their generalizations using object-oriented structuring concepts.

Whilst we now regard recovery blocks, and N -version programming for that matter, simply as special cases of a more general scheme, there has been a somewhat surprising continued interest by others — especially those involved with statistical experiments and with mathematical modeling (for example [Arl90, Puc90, Tai93, Tom93]) — in the basic schemes. This is very flattering, but “real-world” usage of recovery block concepts (see for example [Gil83, Hau85, Gra91, Sim90, Gop91]) has always had to deal with such complexities as input-output, parallelism, hardware faults, etc. — so we would urge more concentration on the richer forms

of structuring for error recovery and for design diversity which have since been developed, and which we have attempted to describe in the later sections of this chapter.

ACKNOWLEDGEMENTS

Our research at Newcastle was originally sponsored by the UK Science and Engineering Research Council and by the Ministry of Defense, but in recent years it has been supported mainly by two successive ESPRIT Basic Research projects on Predictably Dependable Computing Systems (PDCS and PDCS2). Needless to say, as will be obvious from the references we have given, the work we have attempted to summarize here has been contributed to by a large number of colleagues. It would be invidious to name just some of these, but we are very pleased to acknowledge our indebtedness to all of them.

REFERENCES

- [Agh92] G. Agha, S. Frolund, R. Panwar and D. Sturman. A linguistic framework for dynamic composition of dependability protocols. In *Proc. Third International Working Conference on Dependable Computing for Critical Applications*, pages 197–207, Mondello, 1992.
- [Amm87] P.E. Ammann and J.C. Knight. Data diversity: an approach to software fault tolerance. In *Proc. Seventeenth International Symposium on Fault-Tolerant Computing*, pages 122–126, Pittsburgh, 1987.
- [Amm88] P.E. Ammann and J.C. Knight. Data diversity: an approach to software fault tolerance. *IEEE Transactions on Computers*, 37(4):418–425, 1988.
- [Anc90] M. Ancona, G. Doderio, V. Gianuzzi, A. Clematis and E.B. Fernandez. A system architecture for fault tolerance in concurrent software. *IEEE Computer*, 23(10):23–32, 1990.
- [And85] T. Anderson, P.A. Barrett, D.N. Halliwell and M.R. Moulding. Software fault tolerance: an evaluation. *IEEE Transactions on Software Engineering*, 11(12):1502–1510, 1985.
- [And81] T. Anderson and P.A. Lee. *Fault Tolerance: Principles and Practice*. Prentice-Hall, 1981.
- [And76] T. Anderson and R. Kerr. Recovery blocks in action: a system supporting high reliability. In *Proc. Second International Conference on Software Engineering*, pages 447–457, San Francisco, 1976.
- [And83] T. Anderson and J.C. Knight. A framework for software fault tolerance in real-time systems. *IEEE Transactions on Software Engineering*, 9(3):355–364, 1983.
- [Arl90] J. Arlat, K. Kanoun and J.C. Laprie. Dependability modeling and evaluation of software fault tolerant systems. *IEEE Transactions on Computers*, 39(4):504–513, 1990.
- [Avi77] A. Avizienis and L. Chen. On the implementation of N-version programming for software fault-tolerance during execution. In *Proc. International Conference on Computer Software and Applications*, pages 149–155, New York, 1977.
- [Bar83] G. Barigazzi and L. Strigini. Application-transparent setting of recovery points. In *Proc. Thirteenth International Symposium on Fault-Tolerant Computing*, pages 48–55, Milano, 1983.
- [Bar93] P.A. Barrett and N.A. Speirs. Towards an integrated approach to fault tolerance in Delta-4. *Distributed System Engineering*, (1):59–66, 1993.
- [Ben92] M. Benveniste and V. Issarny. Concurrent programming notations in the object-oriented language Arche. Research Report, 1822, Rennes, France, INRIA, 1992.
- [Bon93] A. Bondavalli, F. DiGiandomenico and J. Xu. Cost-effective and flexible scheme for software fault tolerance. *Computer System Science & Engineering*, (4):234–244, 1993.
- [Bon92] A. Bondavalli and L. Simoncini. Structured software fault tolerance with BSM. In *Proc. Third Workshop on Future Trends of Distributed Computing Systems*, Taipei, 1992.
- [Cam79] R.H. Campbell, K.H. Horton and G.G. Belford. Simulations of a fault-tolerant deadline

- mechanism. In *Proc. Ninth International Symposium on Fault-Tolerant Computing*, pages 95–101, Madison, 1979.
- [Cam86] R.H. Campbell and B. Randell. Error Recovery in Asynchronous Systems. *IEEE Transactions on Software Engineering*, 12(8):811–826, 1986.
- [Car91] G.F. Carpenter and A.M. Tyrrell. Software fault tolerance in concurrent systems: conversation placement using CSP. *Microprocessing and Microprogramming*, (32):373–380, 1991.
- [Chi93] S. Chiba and T. Masuda. Designing an extensible distributed language with a meta-level architecture. In *Proc. Seventh European Conference on Object-Oriented Programming*, pages 482–501, 1993.
- [Cle93] A. Clematis and V. Gianuzzi. Structuring conversation in operation/procedure-oriented programming languages. *Computer Languages*, 18(3):153–168, 1993.
- [Cri82] F. Cristian. Exception handling and software fault tolerance. *IEEE Transactions on Computers*, 31(6):531–540, 1982.
- [Gil83] F.K. Giloth and K.D. Prantzen. Can the reliability of digital telecommunication switching systems be predicted and measured?. In *Proc. Thirteenth International Symposium on Fault-Tolerant Computing*, pages 392–397, Milano, 1983.
- [Gop91] G. Gopal and N.D. Griffith. Software fault tolerance in telecommunications systems. *ACM Operating Systems Review*, 25(2):112–116, 1991.
- [Gra91] J. Gray and D.P. Siewiorek. High-availability computer systems. *IEEE Computer*, 24(9):39–48, 1991.
- [Gre87] S.T. Gregory. *Programming language facilities for backward error recovery in real-time systems*. PhD Dissertation, University of Virginia, Department of Computing Science, 1987.
- [Gre85] S.T. Gregory and J.C. Knight. A new linguistic approach to backward error recovery. In *Proc. Fifteenth International Symposium on Fault-Tolerant Computing*, pages 404–409, Michigan, 1985.
- [Gre89] S.T. Gregory and J.C. Knight. On the provision of backward error recovery in production programming languages. In *Proc. Nineteenth International Symposium on Fault-Tolerant Computing*, pages 506–511, Chicago, 1989.
- [Hau85] G. Haugk, F.M. Lax, R.D. Rover and J.R. Williams. The 5 ESS switching system: maintenance capabilities. *AT&T Technical Journal*, 64(6):1385–1416, 1985.
- [Hec76] H. Hecht. Fault-tolerant software for real-time applications. *ACM Computing Surveys*, 8(4):391–407, 1976.
- [Hec86] H. Hecht and M. Hecht. Software reliability in the system context. *IEEE Transactions on Software Engineering*, 12(1):51–58, 1986.
- [Hec89] M. Hecht, J. Agron and S. Hochhauser. A distributed fault tolerant architecture for nuclear reactor control and safety functions. In *Proc. Real-Time System Symposium*, pages 214–221, Santa Monica, 1989.
- [Hoa78] C.A.R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666–677, 1978.
- [Hor74] J.J. Horning, H.C. Lauer, P.M. Melliar-Smith and B. Randell. A program structure for error detection and recovery. *Lecture Notes in Computer Science*, 16:177–193, 1974.
- [Iss93a] V. Issarny. An exception handling mechanism for parallel object-oriented programming: towards reusable, robust distributed software. *Journal of Object-Oriented Programming*, 6(6):29–40, 1993.
- [Iss93b] V. Issarny. Programming notations for expressing error recovery in a distributed object-oriented language. In *Proc. First Broadcast Open Workshop*, pages 1–19, Newcastle upon Tyne, 1993.
- [Jal84] P. Jalote and R.H. Campbell. Fault tolerance using communicating sequential processes. In *Proc. Fourteenth International Symposium on Fault-Tolerant Computing*, pages 347–352, Florida, 1984.
- [Jal86] P. Jalote and R.H. Campbell. Atomic actions for software fault tolerance using CSP. *IEEE Transactions on Software Engineering*, 12(1):59–68, 1986.
- [Kim78] K.H. Kim. An approach to programmer-transparent coordination of recovering parallel processes and its efficient implementation rules. In *Proc. International Conference on Parallel Processing*, pages 58–68, 1978.
- [Kim82] K.H. Kim. Approaches to mechanization of the conversation scheme based on monitors.

- IEEE Transactions on Software Engineering*, 8(3):189–197, 1982.
- [Kim84] K.H. Kim. Distributed execution of recovery blocks: an approach to uniform treatment of hardware and software faults. In *Proc. Fourth International Conference on Distributed Computing Systems*, pages 526–532, 1984.
- [Kim76] K.H. Kim, D.L. Russell and M.J. Jenson. Language tools for fault-tolerant programming. Tech. Memo. PETP-1, USC, Electronic Sciences Laboratory, 1976.
- [Kim89] K.H. Kim and H.O. Welch. Distributed execution of recovery blocks: an approach for uniform treatment of hardware and software faults in real-time applications. *IEEE Transactions on Computers*, 38(5):626–636, 1989.
- [Kim88a] K.H. Kim and S.M. Yang. An analysis of the performance impacts of lookahead execution in the conversation scheme. In *Proc. Seventh Symposium on Reliable Distributed Systems*, pages 71–81, Columbus, 1988.
- [Kim88b] K.H. Kim and J.C. Yoon. Approaches to implementation of a repairable distributed recovery block scheme. In *Eighteenth International Symposium on Fault-Tolerant Computing*, pages 50–55, Tokyo, 1988.
- [Kim90] K.H. Kim and J.H. You. A highly decentralized implementation model for the programmer-transparent coordination (PTC) scheme for cooperative recovery. In *Proc. Twentieth International Symposium on Fault-Tolerant Computing*, pages 282–289, Newcastle upon Tyne, 1990.
- [Kni85] J.C. Knight, N.G. Leveson and L.D.S. Jean. A large scale experiment in N-version programming. In *Proc. Fifteenth International Symposium on Fault-Tolerant Computing*, pages 135–140, Michigan, 1985.
- [Koo87] R. Koo and S. Toueg. Checkpointing and rollback-recovery for distributed systems. *IEEE Transactions on Software Engineering*, 13(1):23–31, 1987.
- [Lee78] P.A. Lee. A reconsideration of the recovery block scheme. *Computer Journal*, 21(4):306–310, 1978.
- [Lee90] P.A. Lee and T. Anderson. *Fault Tolerance: Principles and Practice*. Springer-Verlag, second edition, 1990.
- [Lee80] P.A. Lee, N. Ghani and K. Heron. A recovery cache for the PDP-11. *IEEE Transactions on Computers*, 29(6):546–549, 1980.
- [Liu92] C. Liu. A general framework for software fault tolerance. In *Proc. IEEE Workshop on Fault-Tolerant Parallel and Distributed Systems*, Amherst, 1992.
- [Mae87] P. Maes. Concepts and experiments in computational reflection. *SIGPLAN Notices*, 22(12):147–155, 1987.
- [Mel77] P.M. Melliar-Smith and B. Randell. Software reliability: the role of programmed exception handling. *SIGPLAN Notices*, 12(3):95–100, 1977.
- [Mer78] P.M. Merlin and B. Randell. State restoration in distributed systems. In *Proc. Eighth International Symposium on Fault-Tolerant Computing*, pages 129–134, Toulouse, 1978.
- [Mey92] B. Meyer. *Eiffel: The Language*. Prentice Hall, 1992.
- [Mou87] M.R. Moulding and P. Barrett. An investigation into the application of software fault tolerance to air traffic control systems: project final report. 1049/TD.6 Version 2, RMCS, 1987.
- [Pow91] D. Powell, editor. *Delta-4: A Generic Architecture for Dependable Distributed Computing*. Springer (Berlin), 1991.
- [Puc90] G. Pucci. On the modeling and testing of recovery block structures. In *Proc. Twentieth International Symposium on Fault-Tolerant Computing*, pages 353–363, Newcastle upon Tyne, 1990.
- [Ram88] P. Ramanathan and K.G. Shin. Checkpointing and rollback recovery in a distributed system using common time base. In *Proc. Seventh Symposium on Reliable Distributed Systems*, pages 13–21, Columbus, 1988.
- [Ran75] B. Randell. System structure for software fault tolerance. *IEEE Transactions on Software Engineering*, 1(2):220–232, 1975.
- [Ran84] B. Randell. Fault tolerance and system structuring. In *Proc. Fourth Jerusalem Conference on Information Technology*, pages 182–191, Jerusalem, 1984.
- [Ran93] B. Randell and J. Xu. Object-oriented software fault tolerance: framework, reuse and design diversity. In *Proc. First Predictably Dependable Computing Systems 2 Open Workshop*, pages 165–184, Toulouse, 1993.

- [Rub93] C.M.F. Rubira-Calsavara and R.J. Stroud. Forward and backward error recovery in C++. In *Proc. First Predictably Dependable Computing Systems 2 Open Workshop*, pages 147–164, Toulouse, 1993.
- [Rus80] D.L. Russell. State restoration in systems of communicating processes. *IEEE Transactions on Software Engineering*, 6(2):183–194, 1980.

- [Rus79] D.L. Russell and M.J. Tiedeman. Multiprocess recovery using conversations. In *Proc. Ninth International Symposium on Fault-Tolerant Computing*, pages 106–109, 1979.
- [Sco84] R.K. Scott, J.W. Gault and D.F. McAllister. Investigating version dependence in fault tolerant software. In *Proc. AGARD Conference Proceedings*, 360, 1984.
- [Sco85] R.K. Scott, J.W. Gault and D.F. McAllister. The consensus recovery block. In *Proc. Total System Reliability Symposium*, pages 74–85, 1985.
- [Sco87] R.K. Scott, J.W. Gault and D.F. McAllister. Fault tolerant software reliability modeling. *IEEE Transactions on Software Engineering*, 13(5):582–592, 1987.
- [Shi84] K.G. Shin and Y. Lee. Evaluation of error recovery blocks used for cooperating processes. *IEEE Transactions on Software Engineering*, 10(6):692–700, 1984.
- [Shr78a] S.K. Shrivastava. Sequential pascal with recovery blocks. *Software — Practice and Experience*, 8:177–185, 1978.
- [Shr79] S.K. Shrivastava. Concurrent pascal with backward error recovery: language features and examples. *Software — Practice and Experience*, 9:1001–1020, 1979.
- [Shr78b] S.K. Shrivastava and A.A. Akinpelu. Fault-tolerant sequential programming using recovery blocks. In *Proc. Eighth International Symposium on Fault-Tolerant Computing*, pages 207, Toulouse, 1978.
- [Shr78c] S.K. Shrivastava and J.-P. Banatre. Reliable resource allocation between unreliable processes. *IEEE Transactions on Software Engineering*, 4(3):230–241, 1978.
- [Shr91] S.K. Shrivastava, G.N. Dixon and G.D. Parrington. An overview of the Arjuna distributed programming system. *IEEE Software*, 8(1):66–73, 1991.
- [Shr93] S.K. Shrivastava, L. V. Mancini and B. Randell. The duality of fault-tolerant system structures. *Software — Practice and Experience*, 23(7):773–798, 1993.
- [Sim90] D. Simon, C. Hourtolle, H. Biondi, J. Bernelas, P. Duverneuil, S. Gallet, P. Vielcanet, S. DeViguerie, F. Gsell and J.N. Chelotti. A software fault tolerance experiment for space applications. In *Proc. Twentieth International Symposium on Fault-Tolerant Computing*, pages 28–35, Newcastle upon Tyne, 1990.
- [Str91] Stroustrup. *The C++ Programming Language*. Addison Wesley, second edition, 1991.
- [Sul90] G.F. Sullivan and G.M. Masson. Using certification trails to achieve software fault tolerance. In *Twentieth International Symposium on Fault-Tolerant Computing*, pages 423–431, Newcastle upon Tyne, 1990.
- [Sul91] G.F. Sullivan and G.M. Masson. Certification trails for data structures. In *Proc. Twenty-First International Symposium on Fault-Tolerant Computing*, pages 240–247, Montreal, 1991.
- [Tai93] A. Tai, A. Avizienis and J. Meyer. Evaluation of fault-tolerant software: a performability modeling approach. In C.E. Landweh, B. Randell, and L. Simoncini, editors, *Dependable Computing for Critical Applications 3*, pages 113–135, 1993. Springer-Verlag.
- [Tom93] L.A. Tomek, J.K. Muppala and K.S. Trivedi. Modeling correlation in software recovery blocks. *IEEE Transactions on Software Engineering*, 19(11):1071–1086, 1993.
- [Tyr92] A.M. Tyrrell and G.F. Carpenter. The specification and design of atomic actions for fault tolerant concurrent software. *Microprocessing and Microprogramming*, 35:363–368, 1992.
- [Tyr86] A.M. Tyrrell and D.J. Holding. Design of reliable software in distributed systems using the conversation scheme. *IEEE Transactions on Software Engineering*, 12(9):921–928, 1986.
- [Wel83] H.O. Welch. Distributed recovery block performance in a real-time control loop. In *Proc. Real-Time Systems Symposium*, pages 268–276, Virginia, 1983.
- [Woo81] W. Wood. A decentralized recovery control protocol. In *Proc. Eleventh International Symposium on Fault-Tolerant Computing*, pages 159–164, 1981.
- [Xu93] J. Xu, A. Bondavalli and F. DiGiandomenico. Software fault tolerance: dynamic combination of dependability and efficiency. Technical Report, 442, University of Newcastle upon Tyne, Computing Science, 1993.
- [Xu94] J. Xu, B. Randell, C.M.F. Rubira-Calsavara and R.J. Stroud. Towards an object-oriented approach to software fault tolerance. In *Proc. IEEE Workshop on Fault-Tolerant Parallel and Distributed Systems*, Texas, June 1994.
- [Yon89] A. Yonezawa and T. Watanabe. An introduction to object-based reflective concurrent computation. *SIGPLAN Notices*, 24(4):50–53, 1989.