FACHBEREICH INFORMATIK

Universität Frankfurt

ATTEMPTO: An experimental fault-tolerant
multiprocessor system

M. Dal Cin, R. Brause; J. Lutz
J.W.Goethe-University of Frankfurt,
West-Germany

E. Dilger; Th. Risse
University of Tübingen, West-Germany

ATTEMPTO:
An experimental fault-tolerant multiprocessor system

M. Dal Cin, R. Brause, J. Lutz
J. W. Goethe University of Frankfurt
West Germany

E. Dilger, Th. Risse
University of Tübingen
West Germany

Abstract

This paper describes the overall hardware and software architecture of a fully decentralized, fault-tolerant system. It provides a single-user multi-tasking computing environment. Currently, the system is intended for use as a test-bed for fault-tolerant computing.

1. Introduction

With the rapid decline in the cost of computer hardware it is now feasible to dedicate a multi-processor system to a single user. Consequently, we felt that it makes sense to exploit the advantages of multiple resources provided by a multi-microprocessor system and to develop a single-user fault-tolerant computing environment. In addition, we felt that conventional architectures of fault-tolerant systems have some serious disadvantages:

- Fault-tolerance mechanisms must be explicitly known and used by the user in his programs. This implies special program changes and impedes third-party software.

- The architectural solution is often very specialized, not modular and, therefore, not portable.

- The hardware and software used do not conform industrial standards.

- Fault-tolerance covers only part of the whole system.

- The system is simply too expensive compared with non-fault-tolerant versions.

In order to overcome this situation we adhered to the following design goals:

Our system has been named ATTEMPTO /1/ (A TesTable Experimental MultiProcessor with fault-TOlerance) and is intended to serve the research team as a test-bed for fault-tolerance mechanisms. The prevailing design goals are:

- The user himself should be able to decide for each application job to what extent it should run in a fault-tolerant environment.

- The mechanisms implementing fault-tolerance should be transparent to the user who sees the system as a multitasking monoprocessor system. Consequently all binary non-fault-tolerant programs must run without changes also in a fault-tolerant mode.

- All fault tolerance mechanisms - such as fault-diagnosis, voting or reconfiguration - should be fully decentralized in order that the systems survive the breakdown of single components.

- The system is to be built from conventional hardware parts. Hence most of its fault tolerance is to be implemented in software.

- The fault tolerance mechanisms must be modular and hardware-independent, thus allowing reconfiguration (for the user or for experiments on fault-tolerance) by adding or exchanging software and hardware components.

- Software necessary for fault tolerance mechanisms must be portable. It must be modular, well-structured and written in a high-level language.

- The operating system and the utilities must conform with standard systems.

-   Fault tolerance must cover the entire system (excluding input and output lines).

These requirements led primarily to the development of a modular, hierarchically structured operating system layer /2/ which provides the fault tolerance services of ATTEMPTO.

After a brief overview of the system (Sec. 2) we will present the overall structure of the operating system layer (Sec. 3) and explain our concept of fault treatment (Sec. 4).

2. System Overview

2.1 Hardware

Single-board computers with dual port RAM were chosen as processing nodes (see Fig. 1). Communication between these nodes is provided by a multi-master-bus. Our approach is, however, also implementable on a multibus system in order to enhance the system's fault tolerance with respect to bus errors. To ensure that none of the processing units damages the user's input data, the user input is directly available to all units (by connecting the user terminal to the serial i/o-port of each board). Via the dual-port RAM's a unique logical communication link is established between each pair of processing nodes. (The memory ports are used WRITE-ONLY on global and READ-ONLY on local addresses. The global base addresses of the memory ports are selected from an EC-Code, to hinder addressing of wrong ports by bit faults with memory, bus lines or bus arbiter as possible sources, cf. Fig. 3).

For each processing unit there is one interrupt line on the communication bus. The sender of a message broadcasts its message over the communication bus to all concerned units, including itself. After transferring the data, it activates the interrupt line dedicated to it. This triggers the read of the message by all units providing an asynchronous, atomic transmission of messages; cf. Fig. 4. The temporal order in which incoming messages are accepted is the same for all processing units. It may, however, be different from the temporal order of their individual arrivals. A message transfer protocol specified in /3/ establishes the base for this synchronisation of the processing units.

2.2 Operating System

ATOS is the node operating system of ATTEMPTO. It is comprised of two parts: The OS-Kernel and the Fault-Tolerance Layer (FTL) which is responsible for implementing the fault tolerance. This layer is transparent to the user and is

programmed in Modula-2 /11/. Its location and connection to the operating system is shown in Fig. 2.
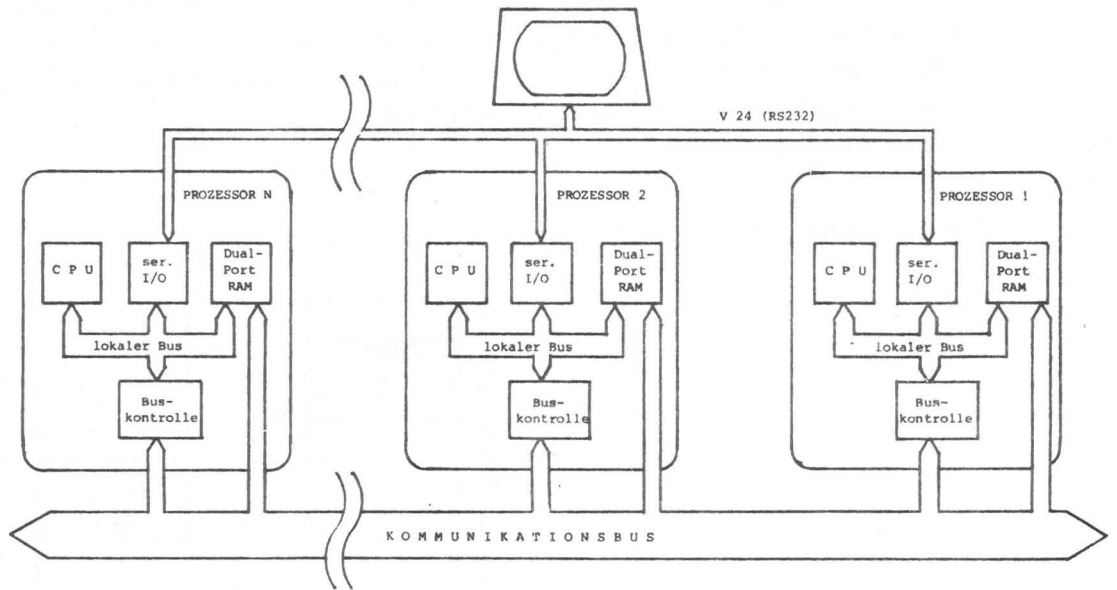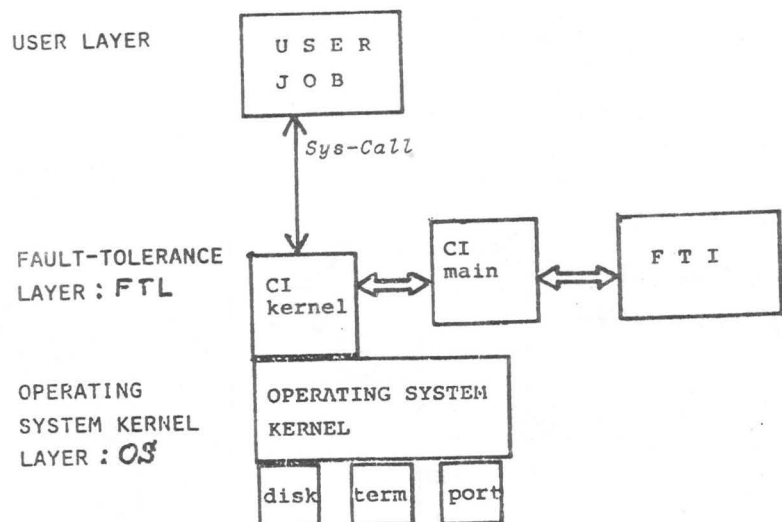
Fig. 1   Hardware

Fig. 2   ATOS

2.3 Job Management

The binding of application jobs to processing units is transparent to the user. It is based on the principle of job attraction /10/ (implemented in software) in order to avoid the need for centralized scheduling and dispatching. Each node maintains its own system tables, which it updates upon receiving messages from other nodes. An (idle) unit applies for the next job by sending a start request to all units including itself. Upon receiving such a request, it marks its system table entry corresponding to the job and responds to the requester. Requests for active jobs (i.e. jobs already being executed by $t+3$ units, see Sec. 4.2) are ignored; cf. Fig. 4.
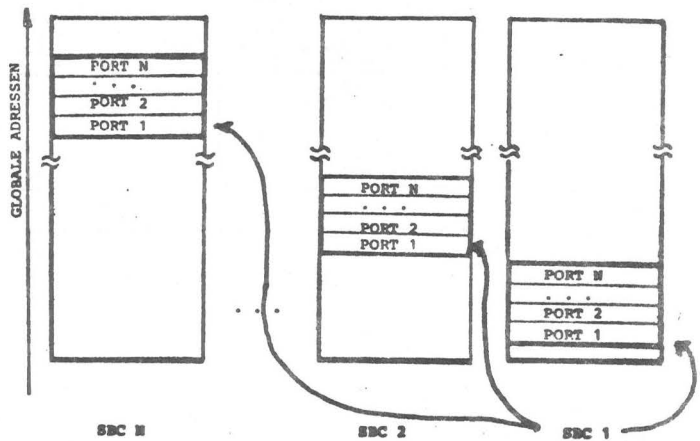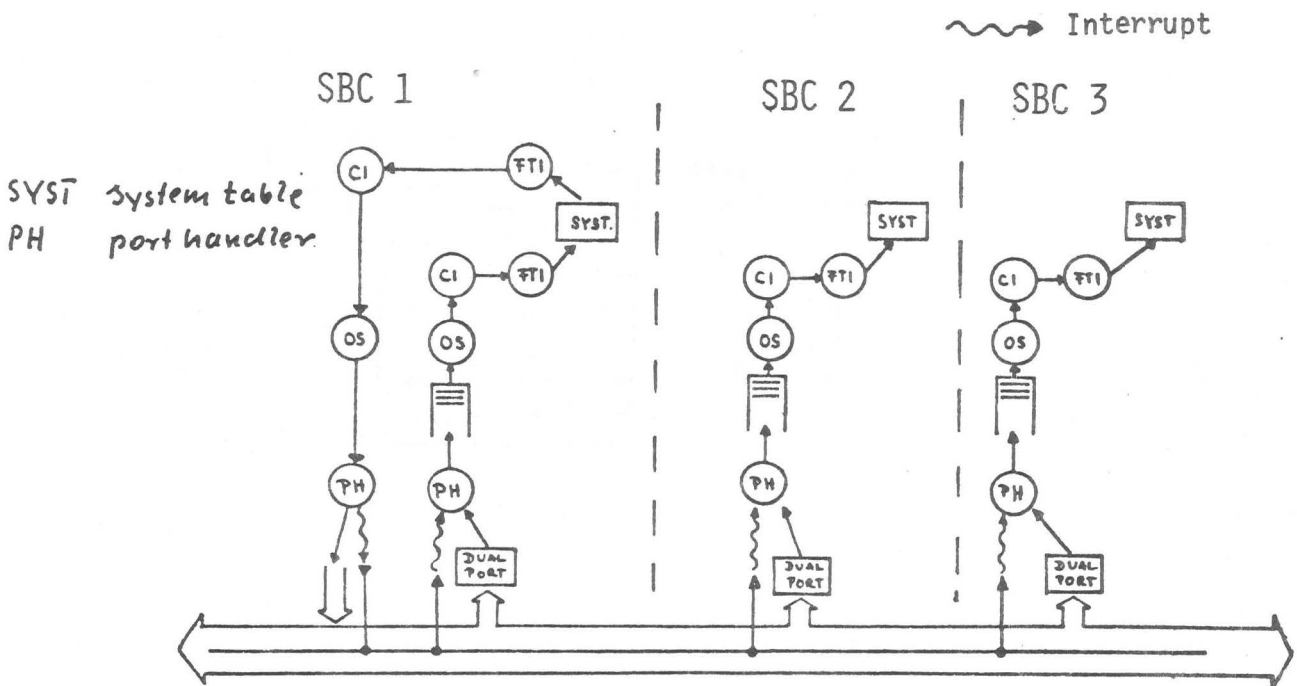
Fig. 3   Address space of message ports

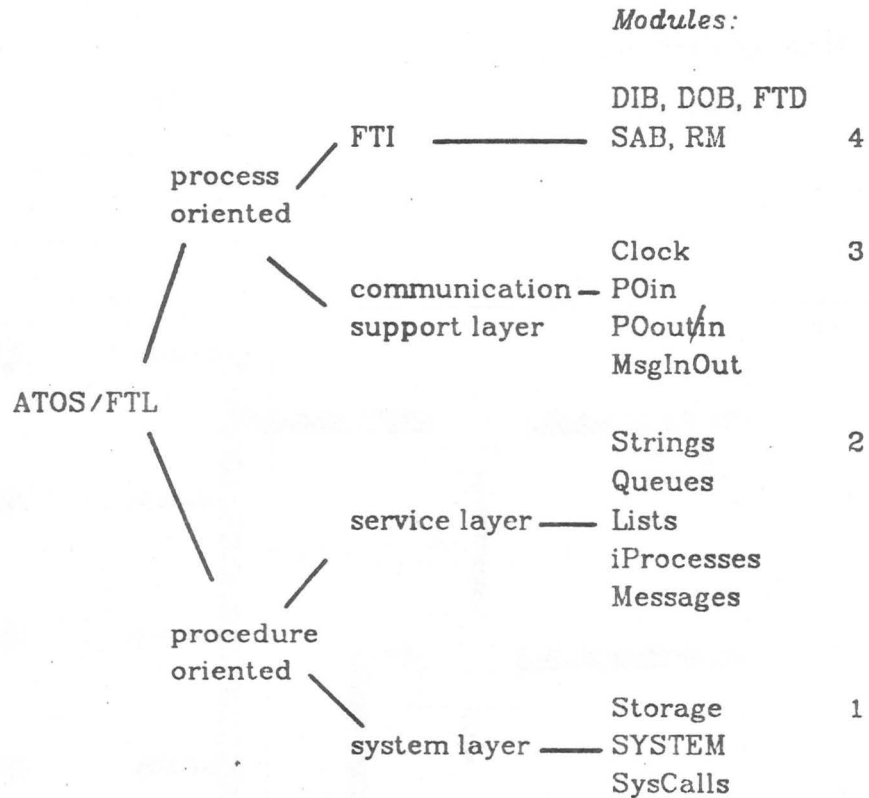Fig. 4   Sending a start-job request

3. The Fault-Tolerance-Layer

3.1 Structure

The FTL of ATOS itself is divided into several functional sublayers /12/ (see Fig. 5), viz:

- the Fault-Tolerance Instance FTI
- the communication support layer
- the service layer and
- the system layer.

Fig. 5   Hierarchy of the modules in the Fault-Tolerance-Layer FTL of ATTEMPTO

Modules:

```
                                                    DIB, DOB, FTD
                                     FTI ────────── SAB, RM        4
                     process   /
                     oriented /
                             /                       Clock         3
                            /          communication ─ POin
                           /           support layer  POout/in
                          /                            MsgInOut
        ATOS/FTL
                          \                            Strings      2
                           \                           Queues
                            \           service layer ─ Lists
                             \         /                iProcesses
                              \       /                 Messages
                            procedure
                            oriented \
                                      \                Storage      1
                                       \ system layer ─ SYSTEM
                                                        SysCalls
```

The sublayers 3 and 4 consist of collections of specific Modula-2 modules. Each module comprises a data strtucture (e.g. Job Control Buffer, Data Input Buffer, Signature Array Buffer, etc.) and an active unit that maintains the data structure. Active units are referred to as Module-clerks and are Modula-2 processes. Clerks communicate by exchanging messages. The second pair of sublayers is composed out of a set of information concealing modules with strictly procedural interfaces.

Hence, the architecture of the higher part of FTL, is based on the message

oriented model of Lauer and Needham /12/ and that of the lower part is based on the procedure oriented model. From the viewpoint of the OS-kernel, the FTL is just another user process (with higher priority) that shares its processor time among several Modula-2 processes.

In ATTEMPTO we must distinguish between three levels of communication:

- communication between clerks (Modula-2 processes)
- inter-process-communication (UNIX processes)
- inter-processor-communication.

Communication between clerks on different SBCs implicitly uses all three levels.

The overall message exchange system of the Modula-2 processes within the UNIX process "FTL" is shown in Fig. 6.

Fig. 6   Message exchange
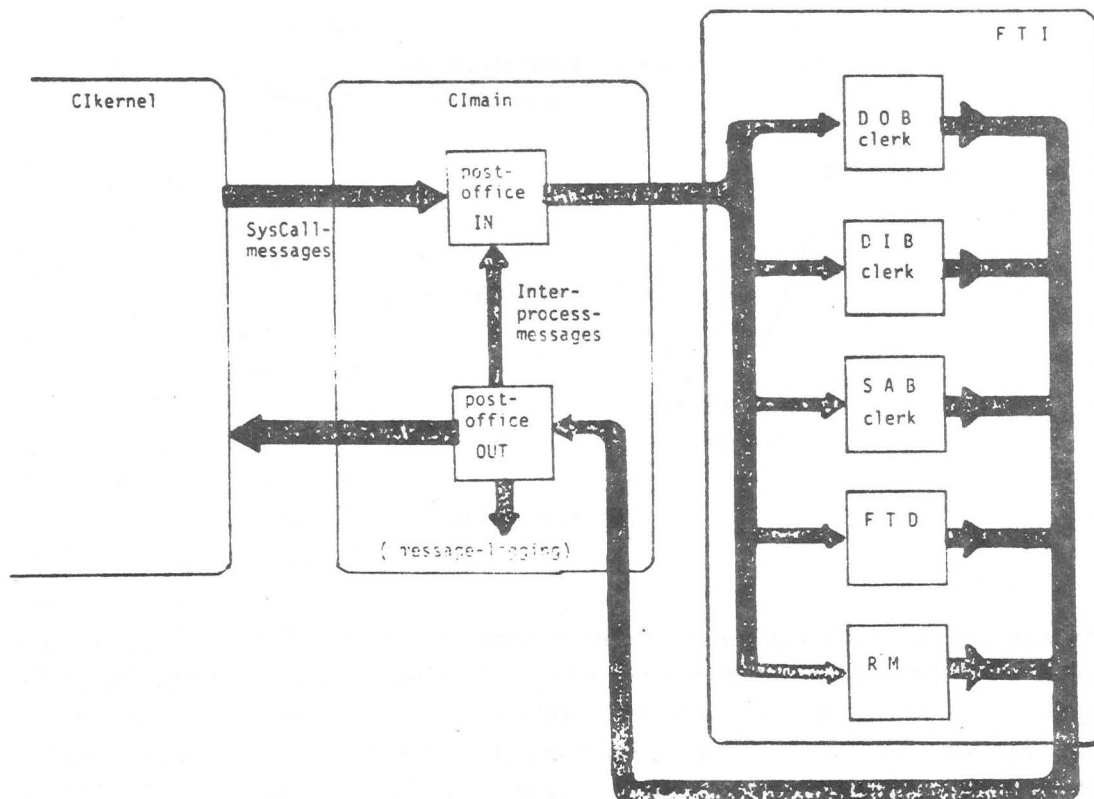
The objective of the FTL is to provide fault tolerance if required by the user. The FTL also provides complete internal observability to the experimenter (but not to the ordinary user). Moreover, its modular and hierarchical structure and the fact that it is exclusively programmed in a higher level language allow us to substitute single modules by modules implementing different strategies for fault-treatment.

3.2 Sublayers

We now characterize very briefly the function of each sublayer: The FTI provides high level fault tolerance services. Its core is formed by the modules DIB, DOB, SAB and FTD. The DIB Clerk manages the data typed in by the user (input buffer) and prepares the user-job output data for fault-diagnosis. The DOB-Clerk manages the user job data output buffer and forwards only the data which are diagnosed as being correct. The DOB-Clerk is authorized to do so by the SAB-Clerk. The SAB-Clerk handles all diagnosis tasks described in Sec. 4. To this end it maintains a so-called signature-array buffer. The FTD-Clerk (Fault-Tolerance Dispatcher) manages a Job Control Buffer and implements the principle of job attraction.

The communication support layer (Communication Instance, CI) is responsible for correct communication between an application job and its FTI as well as between the FTI's of different processing nodes. CI contains a module called Post Office (PO) which constitutes the interface to the OS-kernel. In order to send a message to another node a clerk sends this message to its Post Office. The PO-Clerk completes the message with additional information (e.g. the node-id) and delivers it via the OS-Kernel to the communication port handler. The PO-Clerk forwards also all incoming messages to the receiver clerks of the FTI.

The service layer provides services necessary for Modula-2-process management, buffer management, resource management, etc. The system level provides services for storage management, messages, mailboxes, context switching and system calls.

It is worth mentioning that, although our technique is not specific to any particular implementation of the OS-kernel (it is only essential that the kernel is able to distinguish fault tolerance requests and local system calls), our prototype is intended to run under local UNIX-kernels. Roughly speaking, system calls are diverted to the FTL if fault tolerance requires this. The decision is made by the kernel routine CI-kernel which gains control again as soon as the fault tolerance service has been delivered by the FTL. This technique offers several advantages:

- Every runnable code can be executed fault-tolerantly without modifications in response to the user's wishes (cf. Sec. 4).

- Changes of the kernel that become necessary remain local and controllable since there is only one entry point into the kernel.
- The entry to the FTL is protected just as entries to the OS-kernel are.
- The kernel routine, CI-kernel, can easily be attached to any operating system kernel (pseudo device).
- The method is more or less machine independent.

During development we emulated the system on a minicomputer. Currently the emulation is being upgraded in order to serve as a test-bed for other fault-tolerance purposes and an implementation using single-board-computers with Motorola 680XX and UNIX is under development.

4. Fault Diagnosis and Treatment

With regard to fault-treatment we adhere to an end-to-end strategy /4/. That is, the algorithms which implement fault tolerance are triggerd not before the user-job charges the OS with a WRITE operation. Copies of an application (user) job are executed asynchroneously in parallel by several processing nodes and fault diagnosis is based on the so-called job-result comparison approach /5/. Idle nodes perform self-test routines. This approach is conceptually simple and independent of the hardware structure and of failure types. Several distributed diagnosis protocols for job-result comparison have been investigated and verified by Time-Petri-Net analysis /6/.

Nodes executing copies of an application job form a single virtual processing node. The size of the virtual node is related to the so called degree t of fault tolerance (for the user's job) defined as the maximal number of node breack-downs which can be tolerated (in ATTEMPTO we have t = size of virtual mode - 3 for size > 3, t = 1 for size = 3 and t = 0 else). The degree of fault tolerance can be specified by the user at program start. E.g., typing in "MYPROGRAM#2#" means that "MYPROGRAM" should run with fault tolerance degree 2.

4.1 ATTEMPTO's Diagnostic Model

The classical Preparata-Metze-Chien-Model /7/ (PMC-Model) served us as the basis for the study of diagnostic methods that may be suitable for ATTEMPTO. This model is based on the idea that a system can be partitioned into subunits which test each other. In this case, a test consists of the transmission of a stimulus and of observing the reaction from this stimulus. It is implicitly assumed, that these tests are complete, i.e., that faulty units always show wrong reactions to test stimuli. However, due to the predetermined test direction and the assumption that tests must be complete, this model was abandoned for the use

in ATTEMPTO. The following considerations played a role in our decision:
It is not necessary for the users of a fault-tolerant system that the system is always functioning correctly. Important for the user is that the answers he receives from the system are correct. Hence, not all errors of the system must be treated immediately, rather, just those which make themselves apparent in contact with the environment (end-to-end-strategy).

Therefore, a new diagnostic model based on comparison tests was developed. It forms the basis of the diagnostic procedure used in ATTEMPTO. Comparison tests are employed before any output (WRITE-) operation.

In the PMC-model it also is implicitly assumed that a reliable subunit exists - the so-called "golden unit" - which decides on the basis of the test results which of the subunits are faulty. This assumption rarely applies in real systems. Therefore, we substituted the central diagnosis model of /7/ by a decentralized one.

4.2 Decentralized Diagnosis

The diagnosis of ATTEMPTO begins with the selection of pairs of subunits for comparison. Subunits perform the following steps:
Two units which are spezified by the diagnosis algorithm   compute the same algorithm (user or test program). The respective results are subsequently exchanged and compared. If both results are identical, then both units are assumed to be correct. Rather than to compare all subunits the smallest possible number of pairs is considered. To this end, all possible test assignments have been modeled and analyzed by (undirected) diagnosis graphs similar to the graphs of the PMC model. For ATTEMPTO strictly t-diagnosable, t-optimal diagnosis graphs were chosen /8/. These graphs are the basis of the distributed diagnosis in ATTEMPTO and, for a maximum of t faulty units, are optimal regarding the number of comparisons. (Recall that t is given by the user).

Figure 7 shows two t-optimal graphs with t=4, N=7, and t=5, N=8, respectively. (N number of units executing identical jobs).

Fig. 7   Optimal diagnosis graphs
          ■---■ comparison pair

Every unit sends its results to its neighbors (in accordance with the choosen graph), receives its neighbors results and compares them with its own results. Since the diagnosis graphs are strictly t-optimal, there are at least two units which are neighbors and which can immediately identify themselves as fault-free provided that altogether not more than t units are faulty. All units with erroneous results are neighbors of at least one of these fault-free units. Consequently the fault-free units recognize all faulty ones. A unit with an erroneous result will not find a neighbor with the same result. Nevertheless it may consider itself fault-free.

In order to hinder a faulty unit from passing on its result to the user, further message exchange is necessary. Each unit requires a key (e.g. the initial address of the output routine) in order to output. This key must be sent to it by another unit. (More precisely, unit i asks unit j, for the key. Then unit j returns a message in which the desired key is encoded such that unit i can find the key only if its result coincides with that of unit j. Hence, a faulty unit will not be able to find the key. Recall, that two units are assumed to be faultfree if they produce identical results). The unit which receives the key first is allowed to output. This output is monitored and compared by the other units. In order to limit the bus traffic in ATTEMPTO, the results are compressed to a normal length before they are sent and compared. For data compression a software version of a linear feedback shift register is used as follows:

A data package of length k whose bits are interpreted as coefficients of a polynomial of degree k-1 is divided by a given polynomial of degree r (r=16 in our system) with at least two coefficients $\neq$ 0. The remainder of this division is the signature. Two data packages which differ by one bit produce different signatures /9/. Therefore, all one-bit-faults are detectable. If we, furthermore, assume that all possible faults in a data package are equally possible, we obtain a very low probability P that correct and faulty data packages are not distinguishable by their signatures, viz:

$$P = \frac{2^{k-r}-1}{2^k-1} \approx 10^{-5}.$$

As it can be seen, this probability becomes independent of k for large k. Therefore, it is reasonable to compare large blocks of output data rather than bits or words. This decreases the bus traffic and substantiates our diagnosis assumption that no two faulty units compute the same signatures.

5. Conclusion

In ATTEMPTO the algorithms which implement fault tolerance are triggered, each time a user-job charges the operating system with a WRITE-operation. Consequently, faults of individual subunits are ignored as long as outputs are not produced. Faults are diagnosed and masked using comparison tests just before they become noticeable by a false output or even by a missing one.

Designing ATTEMPTO we confined ourselves to considering only those fault-tolerant concepts which we felt to be fundamental and which did not require extensive hardware modifications. We are, however, convinced that the proposed combination of asynchroneous fault-masking with distributed fault-diagnosis compares favorably with techniques /13/ such as checkpointing and rollback.

References

/1/   Ammann E., Brause R., Dal Cin M., Dilger E., Lutz J., Risse T.: ATTEMPTO, a fault-tolerant multiprocessor work station: design and concepts, Proc. FTCS-13, Milano, pp. 10-13 (1983)

/2/   Risse T., Brause R., Dal Cin M., Dilger E., Lutz J.: Entwurf und Struktur einer Betriebssystemschicht zur Implementierung von Fehlertoleranz, Informatik-Fachberichte 84, Springer, pp. 66-76 (1984)

/3/   Brause R., Ammann E., Dal Cin M., Dilger E., Lutz J.: Softwarekonzepte des fehlertoleranten Arbeitsplatzrechners ATTEMPTO, Symp. German Chapter of ACM, Microcomputing II, Teubner Stuttgart, pp. 328-341 (1963)

/4/   Saltzer J.H. et al.: End-to-end arguments in system design, Int. Conf. Distributed Computing Systems, Paris, pp. 509-512 (1981)

/5/   Ammann E., Dal Cin M.: Efficient algorithms for comparison-based self-diagnosis, Proc. Self-Diagnosis and Fault Tolerance; Dal Cin M., Dilger E. (Eds.): ATTEMPTO Verlag Tübingen, pp. 1-18 (1981)

/6/   Dal Cin M., Florian F. H.: Analysis of a fault-tolerant distributed diagnosis algorithm, Proc. FTCS-15, Ann Arbor, pp. 159-165 (1985)

/7/   Preparata F.P., Metze G., Chien R. T: On the connection assignment of diagnosable systems, IEEE Trans. Electron. Comp. EC-16, pp. 848-854 (1967)

/8/ Ammann E.: Vergleichstestmodelle für selbstdiagnostizierbare Systeme, Informatik Fachberichte 54, Springer, pp. 74-87 (1982)

/9/ Smith J. E.: Measurements of the effectiveness of fault signature analysis, IEEE Trans. Comp. C-29, pp. 510-514 (1980).

/10/ Katsuki D. et al.: PLURIBUS - an operational fault-tolerant multiprocessor, Proceedings IEEE Vol. 66, pp. 1146-1159 (1978)

/11/ Wirth N.: Programming in Modula-2, Springer (1982)

/12/ Lauer H.C., Needham R.M.: On the duality of operating system structures, Operating Systems Review 13, pp. 3-19 (1979)

/13/ Randell B.: System structure for software fault tolerance, IEEE Trans. on Softw. Eng. SE-1, pp. 220-232 (1975)

bisher erschienen:

1/81	Spaniol, Otto:	Performance evaluation of the "virtual subchannel concept" for satellite link communication
2/81	Spaniol, Otto:	Analysis and performance evaluation of DBMS-model
3/81	Wotschke, Detlef Goldstine, Y. Price, Y.K.:	A pushdown automation or a context-free grammar - which is more economical?
4/81	Wotschke, Detlef Goldstine, Y. Price, Y.K.:	On reducing the number of states in a pda
5/81	Kemp, Rainer:	The average height of planted plane trees with m leaves
1/82	Meyer auf der Heide, F.:	Efficiency of universal parallel computers
2/82	Spaniol, Otto Hunsmann, Norbert Nawrath, Franz:	Leistungsbewertung von Rechnersystemen am Beispiel eines Rechnerverbundsystems
3/82	Wegener, Ingo:	The discrete search problem and the construction of optimal allocations
4/82	Wegener, Ingo:	Best possible asymptotic bounds on the depth of monotone functions in multivalued logic
5/82	Meyer auf der Heide, F.:	Infinite cube-connected cycles
6/82	Wegener, Ingo:	Relating monotone formula size and depth of Boolean functions
7/82	Geihs, Kurt:	Analytische und simulative Betrachtung eines Multiplexors mit time-out
8/82	Meyer auf der Heide, F.:	A polynominal linear search algorithm for the n-dimensional Knapsack problem
9/82	Wegener, Ingo:	On the complexity of discrete search problem with positive switch cost
1/83	Jaschinski, Jörg Wegener, Ingo:	Optimal nonadaptive strategies for the search in detection functions
2/83	Meyer auf der Heide, F.:	Efficient simulations among several models of parallel computers

3/83	Wegener, Ingo:	Optimal decision trees and one-time only branching programs for symmetric Boolean functions
4/83	Wegener, Ingo:	Optimal search with positive switch costdecidability, complexity, problems
5/83	Trum, Peter Wotschke, Detlef:	Descriptional complexity for programs
6/83	Meyer auf der Heide, F.:	Lower time bounds for integer programming with two variables
7/83	Wegener, Ingo:	On the complexity of slice functions
1/84	Linnemann, Volker:	Deterministic processor scheduling with communication costs
2/84	Linnemann, Volker:	Grammar rules: a new formalism for defining data types
3/84	Linnemann, Volker:	On the relationship between context free grammars and hierachical data structures
4/84	Jarke, Matthias Koch, Jürgen:	Query optimization in database systems
5/84	Wegener, Ingo:	On the complexity of branching programs and decision trees for Clique functions
6/84	Meyer auf der Heide, F.:	Lower bounds for solving linear diophantine equations on Random Access Machines
7/84	Meyer auf der Heide, F.:	Lower time bounds for testing the solvability of diophantine equations on several parallel computational methods
8/84	Wegener, Ingo:	Time-space trade-offs for branching programs
9/84	Wegener, Ingo:	An improved complexity hierarchy on the Boolean network complexity and formula size of Boolean functions
10/84	Wegener, Ingo:	The critical complexity of all(monotone) Boolean functions and monotone graph properties
1/85	Kintala, Chandra M.R. Wotschke, Detlef:	Concurrent conciseness of degree, probabilistic, nondeterministic and deterministic finite automata
2/85	Seidl, Helmut:	A quadratic regularity test for non-deleting macro s grammars
3/85	Wegener, Ingo:	More on the complexity of slice functions

4/85	Fellmann, Annemarie:	Optimal algorithms for the multiplication in simply generated local algebras
1/86	Bublitz, Siegfried Schürfeld, Ute Voigt, Bernd Wegener, Ingo:	Properties of complexity measures for prams and wrams
2/86	Kemp, Rainer:	The analysis of an additive weight of random trees
3/86	Brustmann, Bettina Wegener, Ingo:	The complexity of symmetric functions in bounded-depth circuits
4/86	Risse, Thomas:	On the symbolical evaluation of the reliability of systems whose structure function is given by any Boolean expression in its components
5/86	Dal Cin, Mario Brause, Rüdiger Dilger, Elmar Risse, Thomas:	ATTEMPTO: An experimental fault-tolerant multiprocessor system