Automatic Control & Computer Sciences, Vol. 23, No. 1, pp. 53-58 (1989)

ALGORITHMIC AND STRUCTURAL ORGANIZATION OF TEST AND RECOVERY FACILITIES IN A SELF-SYNCHRONOUS RING

V. I. Varshavskii, V. Ya. Volodarskii, V. B. Marakhovskii, L. Ya. Rozenblyum, Yu. S. Tatarinov, and A. V. Yakovlev

Avtomatika i Vychislitel'naya Tekhnika, Vol. 23, No. 1, pp. 61-68, 1989

UDC 681.327:519.713

The authors examine algorithmic and structural aspects of the development of a test and recovery automaton that is an important structural element of adapters for a fault-tolerant ring channel of the system interface of a local-area network. The article describes the functions of the adapter in diagnosing and localizing faults in the channel, as well as organization of self-repair.

In developing a fault-tolerant architecture for LANs with a ring channel, an important issue is that of ensuring a high level of fault tolerance for the channel itself, which can constitute the worst bottleneck in the system. To achieve the necessary reliability figures for a ring, the essential diagnostic foundation can be provided by the use of the principle of self-synchronization in developing channel adapter (CA) circuitry. Papers [1,2] considered problems of structural organization, data exchange protocols, and self-synchronous hardware implementation of adapters for ring channels. The basic functions of an adapter for a fault-tolerant ring channel include not only execution of protocol algorithms, but also fault diagnostics in the process of normal operation, fault localization, and self-recovery. In the CA a module was dedicated to these functions (see Fig. 1 in [2]), called a test and recovery automaton (TRA); this module creates a reliable environment for the protocol automaton (PA), which is basically oriented toward implementation of operations of channel bidding and message reception and transmission in "source" and "recipient" modes. In addition, the PA generates fault signals for the FIFO buffer modules, and a variety of auxiliary signals for the TRA. In this paper, we will consider algorithmic and structural aspects of the design of the TRA, whose primary function is that of dealing with faults that occur on channel lines through disconnection of faulty lines and switching in of good ones, as well as detection and dynamic reporting of faults in the CA itself.

1. BRIEF SPECIFICATION OF REQUIREMENTS ON TEST AND RECOVERY AUTOMATON

The process of channel recovery involves successive implementation of three functions: 1) diagnosis of faulty lines, which involves establishment of a malfunction in channel operation; 2) fault localization, which involves the determination of defective lines from both the message-source and message-recipient end; 3) repair by switching, causing faulty lines to be replaced by good or back-up ones. Below we enumerate the requirements on the execution of the principal TRA functions; these are an extension of the general requirements on channel fault tolerance which were given in [1].

Fault of PA or TRA. The process of message transmission around a ring channel is disrupted not only when a line fault occurs on the link between any two adjacent CA, but also upon failure of the PA of any CA, which is a message source or recipient (either selected or unselected). On the other hand, the requisite level of channel reliability may not be achieved because of the fact that a fault has occurred in the TRA itself of some CA. Indeed, assume that the TRA has detected a channel malfunction, but this malfunction cannot be dealt with because the TRA is faulty. Therefore it is natural to require that the TRA detect faults in the PA and in itself, and that it establish a mode of direct message translation via the CA (without sending them to the PA), thus making it possible for the next CA around the ring to detect a line malfunction and to restore the channel.

Naturally, transfer of the CA to a mode of direct signal translation between adjacent ring segments means disconnection of part of the computing hardware, specifically the computer that constitutes the channel subscriber via the CA in question. This can sharply reduce the computational power of the system, but preserves its over-all viability. Thus, a CA fault can be interpreted as a fault of the associated computing facility.

Localization of malfunction in one of the channel segments. Organization of automatic repair for a ring channel is based on the general principles of design of self-repairing aperiodic devices [3]; as noted in [4], however, implementation of self-repair of physical channels requires the solution of a number of additional problems, stemming in this case from the fact that the ensemble of CA connected to the ring cannot have common switching-controlled diagnostic and localization hardware. No connections other than the channel lines themselves are allowed. Therefore the functions described above are implemented in decentralized fashion, with each of the TRA locally implementing a specified test and recovery protocol. The self-diagnostic properties of aperiodic (self-synchronous) circuit engineering can be utilized only for detection of channel line malfunctions. The situation is worse as regards fault localization, since, depending on the type of line malfunction and the instant it occurs, it may happen that a fault on some ring segment is localized by aperiodic circuit-engineering facilities either only from the transmitting-CA side or only from the receiving-CA side, and thus we encounter the problem of transmitting the information on the faulty line from one CA to another either around the faulty channel segment or around the ring. Both situations involve considerable difficulties, and therefore it is advisable to employ the following algorithmic approach to localization of malfunctioning lines.

The TRA of two adjacent CA on a channel segment, between which a malfunction has been detected, are activated by diagnostic signals and beging to execute, in synchronous mode, an algorithm for adequate localization of malfunctioning lines simultaneously from both ends of the segment, successively setting up code combinations on the lines of the segments, checking the correctness of the setting, and recording malfunctioning lines in the fault flip-flops (FF) of these lines. Switching of the FF of the lines should automatically result in replacement of malfunctioning lines by good ones, e.g., by using sliding switching [4]. After proper operation of the channel segment has been restored, both TRA are disconnected and a new message transmission is attempted.

Synchronization of TRA upon execution of localization algorithm. When a malfunction appears in a ring segment, it will be detected (as a result of stoppage of operation) not only by the pair of CA to which the segment is connected. As the buffer pipeline registers of the CA between the sender and the defective segment become filled with information, these CA will successively stop, and their inactivity over a critical time interval will lead to the generation of malfunction-detection signals on the corresponding segments, since this situation is indistinguishable from a malfunction. In other words, a malfunction that is detected by a number of CA will cause them to shift to execution of the localization algorithm. Above it was stated that this algorithm is implemented as a result of paired interaction of TRA in a synchronous mode. The resultant problem of synchronizing the operation of activated TRA should be solved in decentralized fashion, since it is undesirable to incorporate additional lines for handling synchronization signals into the ring. To deal with this problem, we propose to employ a synchronous/asynchronous approach, based on the following two concepts: the localization algorithm should consist of a number of stages that are initiated asynchronously by changes in the states of the channel lines; each stage should be executed synchronously over a limited number of cycles, such that the synchronization signals of two adjacent interacting TRA cannot diverge appreciably in time; after completion of a stage, the local synchronization-cycle generators should stop.

Activation of all TRA in channel. Paired interactions of TRA on adjacent segments are coordinated as follows. The initiator of the different stages of the distributed-localization algorithm is the TRA of the sender, which sends initiating code combinations around the ring alternatingly in both directions. Each successive stage of execution of the algorithm must be initiated only after the execution of the preceding stage has been completed by all the TRA.

This organization of TRA interaction has two important additional advantages. First, it makes it possible to detect a fault of any TRA using simple resources; a TRA is faulty if it has received a combination that initiates the next stage of the algorithm from one of its neighbors, but has not generated a corresponding combination for its other neighbor over the critical time interval, whose value depends only on the duration of the execution of the stage of the algorithm and is independent of the length of the ring (number of CA in the ring). Second, all the TRA in the ring, except for that of the "master" CA (the sender), perform the same operations under the "management" of the master. If

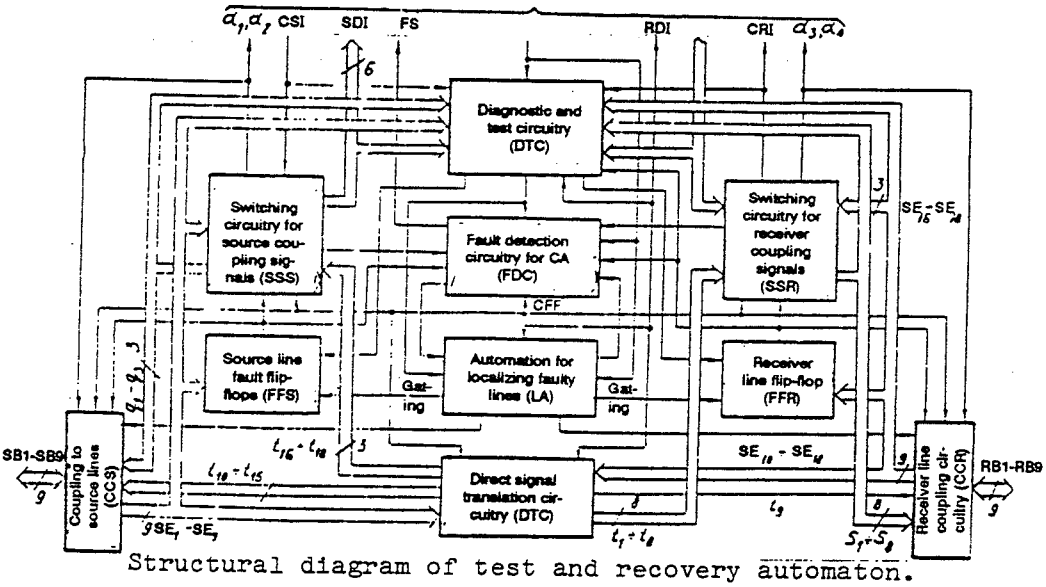

Structural diagram of test and recovery automaton.

it turns out that the master TRA is faulty, it will be eliminated from the ring (its CA will be switched to a direct-translation mode), and execution of the localization algorithm will terminate at the same stage for all the remaining TRA. However, after a new master has been designated by means of systems facilities the localization algorithm will continue to the end, after which there is a return to normal message-transmission operation of the channel.

In order to activate the TRA of all CA around the ring up to and after a malfunctioning segment upon detection of this malfunction, it is sufficient that every connected TRA generate and send around the ring, in the direction of message transmission, a code combination that activates the next successive TRA. In this way, the condition of activation of the TRA of every CA can be either failure to complete the transients on the channel segment from the message-source side over the critical time interval; or the appearance on this segment of a code combination that initiates TRA activation.

Mode of informing computers about faults. If, in the process of message transmissic around the channel, the TRA of some CA detects a fault of its own PA, then this CA must be switched to a mode of direct message translation. This fact naturally leads to disruption of the relative addresses in the ring [1]. Therefore we require that, in this case, all the remaining CA activated their TRA in order to execute the localization and recovery algorithm, after which systems facilities are employed to update the table of correspondence of the absolute and relative addresses of the operting computers connected by the ring. For this, every activation of a TRA must be accompanied by generation of a CA fault signal (FS = 1), which will be reset after the TRA is disconnected. Transfer of a CA to a direct signal translation mode is also accompanied by generation of a signal (FS = 1). A constant FS = 1 value should lead to stoppage of operation of the corresponding computer. A CA can be returned from the direct-translation to the normal mode only by means of an initial-set signal.

2. STRUCTURE OF TEST AND RECOVERY AUTOMATON

In conformity with the above requirements, we developed the structural diagram of the TRA shown in the accompanying figure. It consists of the following: two coupling circuits with the message-source lines SB1-SB9, via which the CA receives messages from its left neighbor, and with message-receiver lines RB1-RB9, via which the CA transmits messages to its right neighbor (CCS and CCR); two switching circuits for message-source and message-recipient signals (SSS, SSR); two sets of source and recipient line fault flip-flops (FFS, FFR); a direct signal translation circuit (STC); a diagnostic and test circuit (DTC); a CA fault detection circuit (FDC); and an automaton for localizing message-source and message-recipient line faults (LA). The double arrows in the figure represent information connections that contain more than one line.

The CCS and CCR consist of bus amplifiers (sensitive elements) and bus transmitters (line drivers), and incorporate control logic for driver operation. These circuits receive data from the channel and transmit it to the channel in normal operating mode, in a mode of execution of the localization algorithm for faulty lines, and in a mode of direct signal translation.

The SSS effects the following connections: input data lines of SDI RA bus to good lines SB1-SB8 of source data (SD) via the bus amplifiers of the CCS; lines of the acknowledgement signal CSI RA to one of the good lines SB7-SB9 of the CS signal via the corresponding CS driver; and output lines t_{16}-t_{18} of the STC to the corresponding lines SB7-SB9 via the bus drivers of the CCS.

The SSS connects the following: the output information lines of the RDI bus of the PA to "good" lines RB1-RB8 of the receiver data (RD) via the CCR bus drivers; the output information lines t_1-t_8 of the STC to the corresponding lines RB1-RB8 via the CCR drivers; and one of the lines RB7-RB9, over which the CR signal is sent, to the line of the CRI signal of the PA via the corresponding bus amplifier.

Both the SSS and the SSR implement the principle of sliding switching [4]. Switching control is implemented by two groups of fault flip-flops: FFS and FFR. The FFS group records faulty lines SB1-SB9, while the FFR group does so for lines RB1-RB9. Information on the number of recorded faults is transferred to the PA by means of signals α_1-α_4 for generating a state word that characterizes the degree of use of back-up lines. The state word of the channel is regularly incorporated into a message packet from the PA to the subscriber computer, for dynamic monitoring of the channel by the systems monitor.

Information on faulty lines of the SB and RB channel is fed to the FFS and FFR from the outputs of elements SE1-SE18 of the CCS and CCR, while information on shorted lines is supplied by the DTC. This information is recorded in the fault flip-flops on the basis of gating pulses issued by the localization automaton LA.

The DTC contains the PA fault diagnostic logic, a circuit that detects and localizes shorts between channel lines, and circuits that detect initiating code combinations on the lines from the source and receiver sides. To generate the corresponding signals, the DTC inputs are fed all the necessary information from outputs SE1-SE18 of the CCS and CCR bus amplifiers, the information outputs of the SDI and q_1-q_3 and s_1-s_8 of the SSS and SSR, the RDI information outputs of the PA, and the inputs of the flip-flops of the FFR group. In addition, the DTC inputs are fed control signals from the PA and LA. The DTC output signals control the operation of the FDC and LA, and send information on shorted lines to the fault flip-flop circuits to be recorded.

The FDC, which is intended to detect faults in the CA itself, incorporates a common fault flip-flop (CFF) together with its activation logic, and a pair of time delay elements (τ_1 = 5 μsec and τ_2 = 100 μsec), that measure two critical intervals. The circuit detects faults of three entities, namely the PA, the channel lines together with the associated TRA hardware, and the TRA itself. To detect PA faults and line malfunctions, a short interval τ_1 is employed, while to detect TRA faults a long interval τ_2 is used, acting as a kind of "watchdog's watchdog." When a PA or TRA fault is detected, the CCF is activated, and it shifts the CA to a mode of direct signal translation (it activates the PA, resets the FFS and FFR, activates the STC, and sets the SSS and SSR to passage of signals from inputs t_{16}-t_{18} and t_1-t_8 to outputs q_1-q_3 and s_1-s_8 respectively).

Upon detection of a channel malfunction, or upon recognition by the DTC of an initiating code combination, the FDC generates an activation signal for the LA, which proceeds to block the operation of the PA and to initiate the localization algorithm. Failure to complete any stage of this algorithm over the critical time interval τ_2 indicates that the TRA is faulty. In this case the FDC activates the CFF.

If when the TRA executes the localization algorithm it turns out that more than two fault flip-flops, belonging to one of the FFS or FFR groups, are activated, this means that it is impossible to restore operation of the corresponding channel segment, because of the lack of back-up lines. In this case, therefore, the FDC also activates the CFF, which halts the operation of the TRA and shifts the CA to a direct-translation mode.

Since, upon activation of the CFF, the CA ceases to execute its main functions, the FDC generates a fault signal FS, by means of which the CA informs its interface communications controller with the subscriber computer. In addition, the FS is generated every time that the LA is activated, and is cancelled after it is switched off, in order to inform the interface controller to initialize the message buffers FIFO1 and FIFO2 and to

initiate operations to update the table of correspondence of absolute and relative computer addresses in the LAN.

After activation, the LA begins to execute the localization algorithm for faulty bus lines, each stage of which (except for the first) is initiated by asynchronous signals from the DTC. Inside each stage of the algorithm, the LA generates a sequence of control signals of fixed length, that establish initiating code combinations on the channel lines (by means of the CCS and CCR bus drivers), provide gating for delivery of information on malfunctioning lines to the fault flip-flops, and check the correctness of switching that results in replacement of faulty lines by good or back-up ones. The LA establishes the duration of the control signals by means of the time delay element τ_1 = 5 μsec of the FDC.

The STC begins to operate when the common-fault flip-flop CFF is activated. It transfers signals from inputs SE1-SE9 to outputs t_1-t_9 and from inputs SE10-SE18 to outputs t_{10}-t_{18}. When the SSS and SSR circuits are appropriately set (flip-flops FFS and FFR reset), this ensures direct signal translation from the lines of the channel segment of the source to the lines of the channel segment of the recipient, and vice versa.

3. ANALYSIS OF CHANNEL FAULTS THAT CAN BE DEALT WITH

Any fault-tolerant system can restore its operational status only provided that the faults that occur are those for which it was designed, i.e., those that it is capable of detecting and localizing. For the fault-tolerant LAN under consideration, therefore, it is necessary to clearly specify the classes of malfunctions with respect to which it will display the property of self-recovery of operation.

The principal classification criteria of faults are the type and multiplicity of the malfunction, and the time and location of its occurrence.

Since we propose to detect channel faults by means of aperiodic circuits, it is necessary to recall the fact, proved theoretically [3], that such circuits belong to the class of completely self-testable circuits with respect to stuck-at conservative faults of any multiplicity at element outputs, or with respect to faults that can be reduced to these. As regards faults in wires and connecting links, these circuits are not completely self-testing, but in this case it is precisely malfunctions of wires (or lines) that must be detected.

We will demonstrate the capacity for detection of faults that can arise on some segment of a baseband channel between a pair of adjacent CA. Simple analysis leads to the following possible malfunctions: 1) stuck-at 0 or 1 on a line; 2) open line; 3) short between any two lines; 4) multiple short between lines. Other types of overly "exotic" faults will not be considered.

The first type of fault is a stuck-at fault and can be referred to the outputs of sensitive elements whose inputs are connected to a malfunctioning line; consequently, they will be detected.

An open line results in establishment of a high voltage level (stuck-at 1) on the part of the line that is not connected to the output of the bus driver. This stuck-at fault can be referred to the output of the corresponding sensitive element, and hence it will also be detected.

Shorting between a pair of lines is not a stuck-at fault that can be reduced either to stuck-at 0 or stuck-at 1, since if two 0's or two 1's are transmitted simultaneously over these lines, a malfunction of this type will not manifest itself at all. However, it is not possible to transmit combinations such as 10 and 01 over shorted lines, since the line over which the 0 is transmitted (low voltage level) will shunt the second line. In this case, detection requires special methods.

We will transmit combinations of self-synchronizing code (SSC) over the lines [1], using a combination consisting of all 0's as a delimiter. Then, if the line over which a 1 should be transmitted is shorted to the line over which a 0 is transmitted, the SSC combination will not be established because of the missing 1's. This situation is easily detected; failure of the SSC combination to appear over the critical time interval indicates that a fault has occurred. In addition, the line on which a 1 should have been set, but was not, can readily be localized on the transmitting CA side.

In the case of another situation, in which an SSC combination is set on the lines, but there is no acknowledgment signal because its line is shorted with a line in the 0 state, this situation will be detected similarly but will be localized only on the receiving CA side. Now, if we isolate the localized line, the second line ought to be usable. However, its electrical characteristics have been disturbed, and there is no guarantee that the line will not be a source of noise. Therefore it is better to isolate both lines. For this purpose, a stuck-at 0 is set on the localized line, and thus a non-stuck-at fault of short-circuit type between lines reduces to a stuck-at fault of multiplicity 2.

Multiple shorting between lines is now of no interest, because it reduces to a stuck-at defect of multiplicity greater than 2, and cannot be combatted because of the lack of spare lines (see the requirements on a channel in [1]).

In terms of the place of occurrence, we should distinguish the following faults: channel-line faults; faults in TRA hardware associated with lines and switched together with them; faults in nonswitching TRA hardware; and PA hardware faults.

Channel-line defects were just considered. Lines are connected to certain TRA hardware that is switched together with them. This hardware includes all sensing elements and drivers together with their control logic, direct-translation circuitry, and a variety of switching elements. Only stuck-at faults at element outputs can be detected in this part of the hardware, since these manifest themselves as stuck-at faults on the lines and hence can be localized and averted together with their lines.

Stuck-at defects at gate outputs in the rest of the TRA hardware and in the PA circuitry can be detected but not corrected, since they have no back-up. Special diagnostic devices are introduced for this purpose, these comprising straightforward diagnostic logic and elements for generating critical time intervals. When such faults are detected, the DFF of the common CA fault should be activated, putting the CA in a direct-signal-translation mode.

In concluding, let us consider the permissible fault multiplicity. The theory of aperiodic circuits asserts that, by means of such circuit-engineering techniques, it is possible to detect stuck-at faults of arbitrary multiplicity at element outputs. However, this assertion discusses only a possibility, which still needs to be implemented. Indeed, if an aperiodic device has only one diagnostic circuit, then a single fault in this circuit will cause a single defect in the device itself to go undetected. Therefore it must be required that the PA and nonswitching TRA hardware detect, as a minimum, single faults of stuck-at type at element outputs. As for the switched TRA hardware and the backbone lines, the given scheme provides for combatting faults of multiplicity not greater than 2 and detection with a multiplicity greater than 2 on every ring-channel segment.

This paper is the concluding one in a series (see [1,2]) dealing with the design of fault-tolerant microcomputer LANs of standard type (with object interface of Q-bus type), based on ring-channel architecture. Our basic arrangement of a self-synchronous channel adapter is currently in the prototypuing stage. In the future, we intend to develop an LSI chip set for this type of channel. Our work promotes an evolutionary strategy toward the deployment of fault-tolerant self-synchronizing circuitry on the physical layer of network architecture [4].

REFERENCES

1. V. I. Varshavskii, V. Ya. Volodarskii, V. B. Marakhovskii, L. Ya. Rozenblyum, Yu. Ya. Tatarinov, and A. V. Yakovlev, "Structural organization and data exchange protocols in a fault-tolerant self-synchronous ring," AVT [Automatic Control and Computer Sciences], no. 4, pp. 48-55, 1988.

2. V. I. Varshavskii, V. Ya. Volodarskii, V. B. Marakhovskii, L. Ya. Rozenblyum, Yu. S. Tatarinov, and A. V. Yakovlev, "Hardware implementation of protocols of a fault-tolerant self-synchronous ring," AVT [Automatic Control and Computer Sciences], no. 6, pp. 60-69, 1988.

3. Automaton Control of Asynchronous Processes in Computers and Digital Systems [in Russian], Nauka, Moscow, 1986.

4. V. I. Varshavskii, V. B. Marakhovskii, L. Ya. Rozenblyum, Yu. S. Tatarinov, and A. V. Yakovlev, "Approach to reliable circuit-engineering implementation of physical layer protocols of network architecture," AVT [Automatic Control and Computer Sciences], no. 6, pp. 76-81, 1986.

12 February 1988