A Microcomputer-Based Fault-Tolerant Control System for Industrial Applications

ASOK RAY, SENIOR MEMBER, IEEE

Abstract—**The concept and development of a fault-tolerant control system and the results of its real-time application using commercial microprocessors are presented. The system incorporates (via recursive filtering) on-line detection and reconfiguration of faulty equipment, sensor calibration, and information display within a structure that relies on alternatively configured regulators for plant control under different operational modes; these regulators are designed on the basis of analytically derived control laws and using a rule-based heuristic approach. Each critical input signal is a weighted average of all valid measurements of the appropriate plant variable where the weighting matrix is adaptively updated as a function of *a posteriori* probabilities of failure of these redundant measurements. Since the weight of a degraded measurement is smoothly reduced, eventual isolation of this fault does not cause an abrupt change in the estimate of the measured variable, and therefore, the feedback control system remains "bumpless." The automated system has been shown to be tolerant of equipment failure(s), sensor degradation and noise, and certain types of process malfunctions, disturbances, and uncertainies by experimentation at the MIT nuclear research reactor.**

INTRODUCTION

WITH THE ADVENT of inexpensive and reliable microprocessors during the last decade, small computers are being adapted as part of instrumentation and control in diverse industrial processes. They play a dominant role in the evolution of fault-tolerant control systems. A fault-tolerant system can be viewed as an integrated system that has inherent capabilities to detect, isolate, and reconfigure failures of equipment (including its own components). Although a number of fault-tolerant control systems have been designed and tested for avionic applications [1]–[3], such systems have not yet attained a similar level of acceptance in other major industries such as chemical and power plants. Recently, Stanley [4] has implemented a closed-loop control system in a chemical plant that makes use of redundant measurements to detect and isolate some of the faulty instruments and to estimate relevant plant parameters. Although Stanley's work is significant in the evolution of fault-tolerant control technology in chemical industries, further research and development efforts are required for its commercial acceptance.

On-line applications of computer for closed-loop control of nuclear power plants has been, excepting the Canadian HWR units [5], quite limited. Recently, several computer-aided fault diagnostic techniques have been developed with the objective of improving operational safety in nuclear power plants [6]–[9]. Apparently, the majority of these techniques are not designed for closed-loop control and other tasks that are directly related to improved plant performance. Nevertheless, it is possible to achieve simultaneously superior safety, reliability, and operational flexibility and performance of complex industrial processes such as chemical and power plants. This can be accomplished by integrating fault detection, isolation, and reconfiguration (FDIR), sensor calibration, measurement estimation, and information display techniques within a given control system structure of complex industrial processes.

A fault-tolerant control system software has been developed and designed to be executable on commercially available microcomputers in an industrial environment. In an ongoing series of experiments, the control system software has been implemented in the 5-MW nuclear research reactor MITR-II using an LSI-11/23 microcomputer system and compatible data acquisition and reconstruction devices. The nuclear reactor MITR-II is licensed by the Nuclear Regulatory Commission and is operated by the Massachusetts Institute of Technology. Presently, the power of MITR-II can be digitally controlled, under steady-state and transient operations, via feedback of a number of sensor signals such as neutron flux, primary coolant flow, temperature, and the regulating rod position. The controller uses heuristics and learning theory for emulating operator instructions. A real-time CRT display presents the validated data and diagnostics of the relevant instrumentation and equipment.

The major functions of the on-line fault-tolerant control system at MITR-II are

1) the regulation of reactor power under steady-state and transient conditions such as xenon oscillations, coolant temperature variations, and rapidly changing load;
2) the on-line detection, isolation, and reconfiguration of faulty sensors without interrupting the plant operation;
3) the on-line estimation of both measurable and non-measurable plant variables such as power, coolant flow, temperature, and reactivity using the available sensors and/or analytic measurements [10] that are obtained from real-time models using physical relationships among the plant variables;
4) the on-line calibration of power sensors to compensate for process disturbances such as changes in the spatial distribution of neutron flux due to xenon transients;

Paper IUSD 83-6, approved by the Industrial Control Committee of the IEEE Industry Applications Society for publication in this TRANSACTIONS. Manuscript released for publication February 4, 1985. This paper was supported by the Charles Stark Draper Laboratory under a research grant to the Massachusetts Institute of Technology.

The author is with the Mechanical Engineering Department, Pennsylvania State University, University Park, PA 16802.

0093-9994/85/0900-1276$01.00

5) an on-line information display of the critical plant variables and diagnostics of faulty sensors and equipment, assisting the operators to make timely and appropriate decisions.

The software for the fault-tolerant control system of MITR-II is structured such that it can be readily adapted and extended to commercial power plants and complex chemical processes. Part of the system software that involves detection and isolation of faulty sensors has already been applied to pressurized water reactor (PWR) [11], boiling water reactor (BWR) [12], and breeder reactor power plants [13].

The objectives of this paper are to present the development and on-line implementation of the fault-tolerant control system and to elucidate its characteristics. The paper is organized in several sections which describe the salient concepts, instrumentation and control system of the test facility, system software, and experimentally deduced results and conclusions. The analytical development of the methodologies used in the software, experimental techniques, and associated safety features have been reported in previous publications which are listed in the references.

CONCEPTS OF FAULT TOLERANCE, CALIBRATION, AND ESTIMATION

The signal validation methodology uses redundant measurements that may be either direct sensor outputs or analytically obtained from a mathematical model formulated on the basis of physical relationships among other process variables (e.g., mass and energy balances in thermofluid processes). It provides a unified systematic procedure for 1) FDIR and 2) sensor calibration and measurement estimation.

The FDIR technique used in the fault-tolerant system is a sequential decisionmaking procedure that systematically seeks out the largest consistent subset from a set of redundant measurements where the consistencies among individual measurements of a given process variable are determined on the basis of allowable errors [8], [9]. Allowable errors or error bounds specific to individual measurements can be obtained from either experimental data or the information on instrument tolerances. As the number of redundant measurements increases, the checking of consistency in all possible combinations and the attendant task of bookkeeping for all information at the current and past sampling instants become very complex. A systematic procedure, appropriate for a digital processor, that relies on recursive relations based on the consistencies of each measurement relative to the remaining ones has been developed for diagnosing the sensor and plant equipment failures. The mathematical background and other details of the FDIR technique are given in the previous publications [8], [9].

The sensor calibration and measurement estimation technique is also a sequential procedure which is performed on-line in the framework of the aforesaid FDIR technique [8], [9]. If the process variables being measured are time-dependent and if the redundant sensors are installed in different spatial locations (e.g., neutron flux detectors) or if analytic measurements [10] are used to supplement the sensor redundancy (e.g., fluid temperature obtained from mass and thermal power balances), the measurements of the given process variable may exhibit deviations from each other after a length of time even though the sensors are functioning normally. These differences could be caused by time-varying plant parameters, reaction kinetics, transport delays, etc. Consequently, some of the sensors may be erroneously deleted unless they are periodically recalibrated. On the other hand, failure to isolate a degraded sensor could cause an inaccurate estimate of the measured variable, and the plant performance may be adversely affected if that estimate is used as an input to the controller. These difficulties can be circumvented if 1) only the consistent measurements are calibrated such that their residuals are minimized, and 2) the weights of the calibrated measurements (for computing the estimate) are updated on the basis of their respective *a posteriori* probabilities of failure instead of being *a priori* fixed. In the event of abrupt disruptions in some sensor(s) in excess of the specified error bound(s), the respective sensor(s) are isolated by the FDIR algorithm, and only the remaining sensors are calibrated and used to provide an estimate. If a gradual sensor degradation occurs, the faulty sensor may not be immediately isolated, but its influence on the estimate and the calibration of the remaining sensors is diminished as a function of its degradation because its weight decreases with an increase in the *a posteriori* probability of failure. Thus, if the error bounds of the measurements are appropriately increased to reduce the probability of false alarms, the resulting delay in detecting a gradually degrading sensor could be tolerated because an undetected fault, as a result of its reduced weight, does not significantly affect the accuracy of calibration and estimation. Moreover, since the weight of a gradually degrading measurement for computing the estimate is smoothly reduced, the eventual isolation of the fault does not cause an abrupt change in the estimate, thereby assuring a bumpless controller action. Theoretical development of the calibration and estimation technique and relevant experimental results have been reported in a recent publication [14].

INSTRUMENTATION AND CONTROL SYSTEM

A description of the system configuration and instrumentation of the 5-MWt fission reactor is given in the *MITR-II Reactor Systems Manual* [15]. The reactor is heavy-water reflected, light-water moderated and cooled, and functions as a research and educational facility. The nuclear instrumentation used for the research reported in this paper consists of three neutron flux sensors and a gamma-ray sensor that correlates neutron power with the radioactivity (N-16) of the primary coolant. Four independent measurements of primary coolant flow are obtained from the pressure differences across orifices. Primary coolant temperatures are measured as follows: two sensors for the hot leg, two sensors for the cold leg, and one sensor for temperature difference between the legs. The noise and statistical characteristics of the MITR-II's flow, temperature, and neutron flux instrumentation are similar to those in commercial reactors. These sensors are hard-wired to a portable LSI-11/23 microcomputer system

through appropriate isolators, signal conditioners, and A/D converters.

Coarse control of the power in the MITR-II is achieved by manually positioning a bank of six shim blades. Once critical, the neutron flux is normally maintained constant by an automatic analog controller that is monitored by the reactor operator. The controller adjusts a fine-control regulating rod according to feedback signal of a single power sensor. The analog controller is a standard proportional-integral-derivative (P-I-D) type. The controller output energizes a three-position relay to drive a constant-speed motor that moves the regulating rod up or down at 108 mm/min, which corresponds to an average of reactivity change of about $1.1 \times 10^{-5}\ \Delta K/K$ per second.

The digital control system shown in Fig. 1 replaces the single power sensor in the analog controller by the four available power sensors, substitutes the analog controller and its relay logic by a digital algorithm, and incorporates on-line fault diagnosis, sensor calibration, and information display. The constant-speed motor has also been replaced by a variable-speed stepping motor in order to implement continuous-action control laws. (*Note:* The maximum speed of the motor is limited to reduce the consequences of any unforeseen catastrophic malfunction.)

The multiply redundant sensors for power, flow, and temperature difference measurements are first validated on-line with the aid of a fault detection and isolation (FDI) methodology. The MITR-II's flow and temperature sensors are stable and do not require on-line calibrations. However, the neutron power sensors may be affected by process disturbances. This occurs because these sensors, which measure the leakage flux from the core, are spatially fixed. Any process that disturbs the leakage flux will cause the output of the neutron sensors to vary even though there may have been no net change in the overall reactor power. For example, as short-lived fission product poisons build into the core, the shim bank must be withdrawn to provide compensating reactivity. This sequence entails no change in reactor power but, as the shim bank is withdrawn, the magnitude of the peak in the axial neutron flux profile decreases and its position relative to the bottom of the core rises. Hence the outputs of the neutron power sensors may change even though reactor power, as indicated by heat balances, is constant.

Given that process disturbances such as xenon transients, coolant temperature variations, and movement of individual shim blades can affect the neutron power sensors and since the reactor control system feeds back the validated power estimate at each sampling interval, the neutron power sensors are calibrated on-line, and the estimate is obtained as a weighted average of all valid calibrated power measurements. The flowchart for the calibration filter is shown in Fig. 2. The output of the calibration filter is protected against possible drifting by comparing it with an analytic measurement that is valid for both steady-state and transient operations. A real-time thermal–hydraulic model of the primary coolant system that accepts the validated measurements of power and coolant flow estimates as inputs is used to generate an analytically redundant value ΔTa for the hot-leg to cold-leg temperature

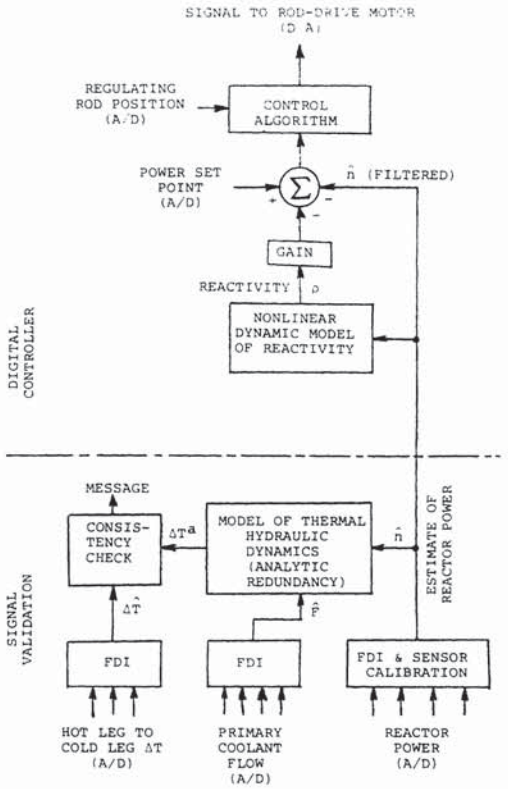

Fig. 1. Simplified digital control scheme for nuclear reactor.

difference [10]. The value of ΔTa is compared with an estimated value of the temperature difference ΔTe obtained from the sensor outputs at every sample. An inconsistency of ΔTa and ΔTe implies, among other potential malfunctions, possible errors in the power estimate since it is one of the inputs to the analytical relationship for computing ΔTa. If this inconsistency occurs, an alarm message is generated to notify the reactor operator. This procedure also guards against common-mode failures, i.e., simultaneous and identical failures of all redundant sensors, possibly due to a common cause, that cannot usually be detected from either power or ΔT sensors alone.

In addition to feeding back the power estimate, the controller is routinely calculating and feeding back the reactivity deviations. An inverse kinetics algorithm is used for on-line evaluation of the reactivity from the reactor's power level and history [16]. The advantage of reactivity feedback is that whenever the equilibrium critical condition is disturbed, the reactivity promptly changes, thereby creating an appropriate control signal that will adjust the regulating rod's position to restore the original condition. Reactivity feedback provides a fast anticipatory control action that is independent of the steady-state value of the reactor power.

The regulating rod's associated reactivity ρ is a nonlinear function of the rod position x. The reactor power is regulated by controlling the reactivity which, in turn, is influenced by the regulating rod's position. Therefore, overall plant gain is directly related to $d\rho/dx$, which varies nonlinearly with changes in the regulating rod's position. In order to compensate for the nonlinear relationship between rod position and

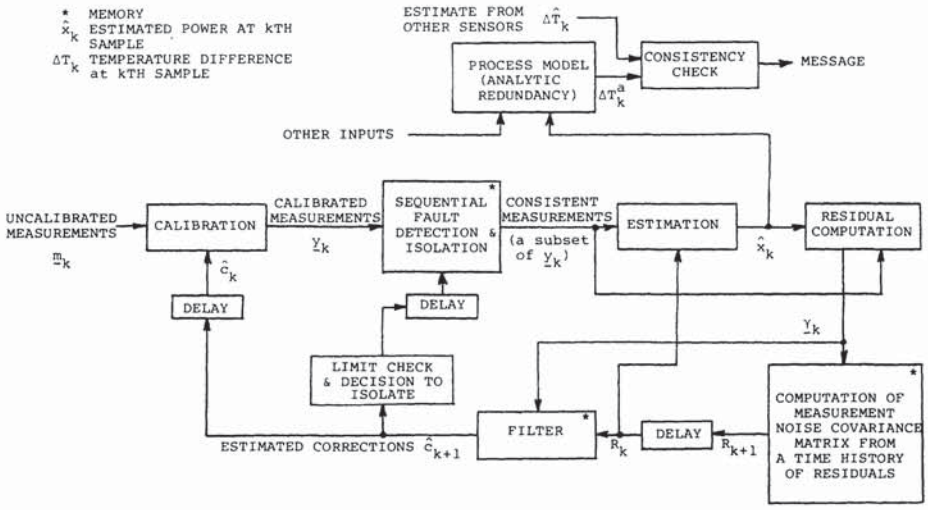

Fig. 2. Flowchart for fault detection and calibration filter.

reactivity, the controller gain is constrained to be proportional to $dx/d\rho$. This is accomplished by approximating the regulating rod's differential reactivity work curve as a piecewise linear function and by measuring the rod's height at each sampling interval.

A nonlinear digital control law was developed using a point-kinetic model which is valid around the critical operational modes of the reactor and is suitable for normal steady-state operations and small perturbances. For large deviations in power such as those that may occur if the reactor is required to follow rapid load changes, some other approach has to be taken for controller design because the reactor physics under these circumstances is governed by very complex equations and control systems designed on the basis of plant model(s) would not be practical. Since nuclear reactor control is currently the domain of human operators who rely on written instructions and past observations, a heuristic control law which tends to emulate human decisions was adopted. The heuristic control algorithm is based on the principles of adaptive and learning theory [17], [18] using a recursive relationship and was implemented in the LSI-11/23 microcomputer as part of the fault-tolerant control software. The control law consists of

- four heuristics defining four different modes of control rod motion;
- five reactor states—one ideal state and four nonideal states defined in terms of reactor period and percent power deviation;
- a penalty function for counterproductive actions; and
- a reward function for the time interval during which the reactor remains in the ideal state following a given control action.

The heuristic controller, although far from being optimal, is much more safe and dependable than a model-based controller for rapid transients. Therefore, under steady-state conditions when the reactor is in the ideal state for a prolonged period, the strategy is to transfer automatically to the model-based nonlinear control law. Upon initiation of a disturbance, if the reactor moves to a nonideal state, the strategy is to transfer immediately to the heuristic control law. An elaborate description and mathematical details of the heuristic controller are reported in a recent publication [19].

SYSTEM SOFTWARE

The fault-tolerant control system software is implemented on an LSI-11/23 microcomputer with KEF-11 floating point hardware, real-time clock, and A/D and D/A conversion capabilities. The software is designed to follow a top-down hierarchical structure as shown in Fig. 3. It is written in standard Fortran and Macro-11 assembler language and is supported by the RT-11 operating system of the Digital Equipment Corporation. The software is specially designed so that it can be easily translated into a structured language such as Pascal and "C." A machine executable form of the program requires a memory of about 30 kbytes that include the libraries of Fortran and special real-time subroutines. The execution time is less than 190 ms/cycle if no messages are generated.

The software system is designed to have a main program, a single data file, and a number of modular subprograms that are dedicated to specific tasks (see Fig. 3). A subprogram may call any number of lower level subprograms (including none). To facilitate software management, each subprogram (but not the data file) is called by exactly one supervisor program and is restricted to communicate with its immediate supervisor and subordinate modules only. Presently, the program is executed sequentially on a single processor. The software structure will allow systematic parallel operations of individual tasks (corresponding to respective subprograms) if the existing single processor is replaced by an appropriate distributed processor. Following a bottom-up approach, a low-level subprogram (such as the decision and estimator subprogram DECEST in Fig. 3) can be replaced by functionally equivalent hardware. Only its supervisor module needs to be updated to interface with the hardware, while the rest of the program remains unaltered.

The main program, called MAIN, acts as a command interpreter and calls the nonreal-time and real-time subpro-

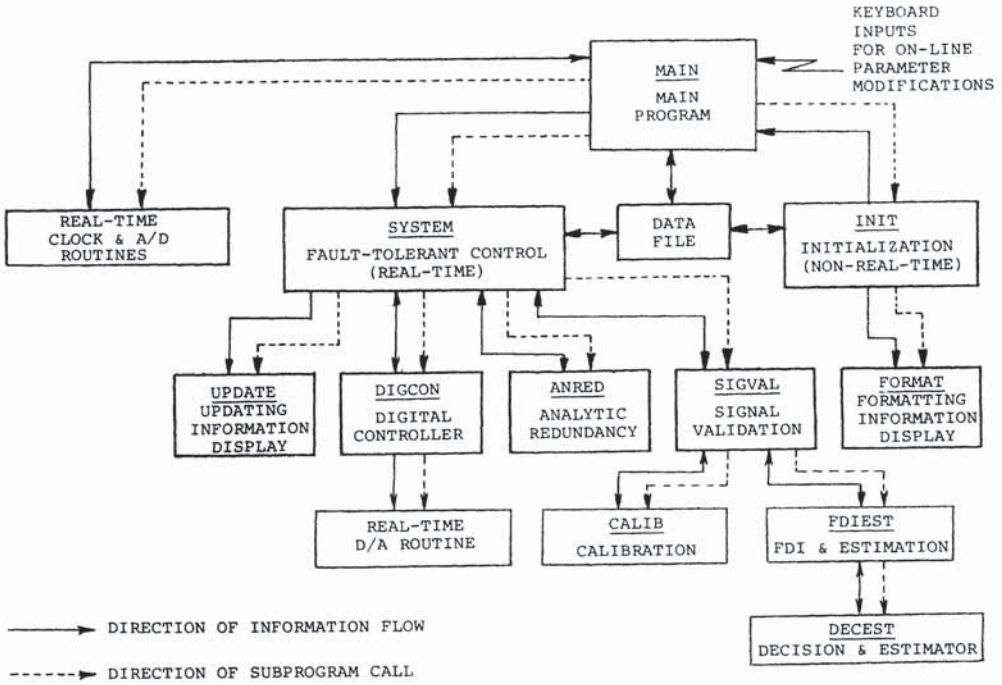

Fig. 3. Hierarchical structure of system software.

grams (refer to Fig. 3). The nonreal-time subprogram INIT is intended for initialization of system parameters, selection of operating modes, and formatting the information display structure (via the subprogram FORMAT) on CRT for man–machine interface. At the end of the initialization phase, MAIN waits to receive a command from the human operator to enter the real-time mode. At the onset of the real-time mode, MAIN cyclically calls the clock and A/D routines, and the fault-tolerant control system subprogram, SYSTEM, at a sampling frequency which was determined during the nonreal-time phase. During the course of operation, certain important decisions such as deletion/insertion of critical sensors and updating of controller parameters can be carried out on-line by keyboard inputs. MAIN interprets these commands with a priority higher than its normal execution.

The real-time subprogram SYSTEM receives the sensor signals from MAIN and then supervises the execution of several tasks. SYSTEM also updates the displayed information periodically via the subprogram UPDATE, where the frequency of information updating is controlled by MAIN and can be altered on-line by a keyboard input. As described in the instrumentation and control section, the tasks performed during each sample are executed by the following subprograms (see Fig. 3).

1) Signal validation subprogram SIGVAL receives a set of measurements for each process variable, detects and isolates failures, if any, and then finds an estimate of the measured variable from the valid measurements only. This subtask is performed by the FDI and estimation subprogram FDIEST which, in turn, calls the lower level decision and estimator subprogram DECEST. During each cycle, FDIEST is called three times, once for each process variable, namely power, flow, and temperature difference. After control is returned to the calling subprogram SIGVAL, subprogram CALIB may be called to calibrate the consistent measurements of the given process variable. As explained earlier, only the power sensors need to be calibrated on-line.

2) The analytic redundancy subprogram ANRED computes the value of a specified process variable (temperature difference in this case) from the valid estimates of other process variables (power and flow). It is a real-time model that relies on the physical laws of dynamic energy and mass balance.

3) The digital controller subprogram DIGCON receives the estimate of power which was transmitted to SYSTEM by SIGVAL. The estimate is used to evaluate reactivity and to generate an actuator command signal via a given control algorithm such as the heuristic control law. The actuator command signal is received by the D/A module, and the resulting analog output is electronically processed to drive the regulating rod.

To enhance the reliability of the software, a systematic test procedure has been developed. Whenever any modifications are made in the software, which may range from simply updating a parameter in the data file to reformulation of one or more subprograms, the resulting load module is systematically tested. Backup copies of immediately earlier versions of the source, compilation list, link list, and load module files are retained, and these modifications are recorded in a log book. The software testing procedure broadly consists of three phases: 1) creation of an executable load module that is free of any error or warning messages; 2) execution of the nonreal-time part of the program and checking the displayed values of critical parameters before entering the real-time mode; and 3) real-time execution of the program under manual control involving tests for fault diagnosis, sensor calibration, and sensitivity of the D/A converter under induced disturbances, and finally operating the plant under automatic control.

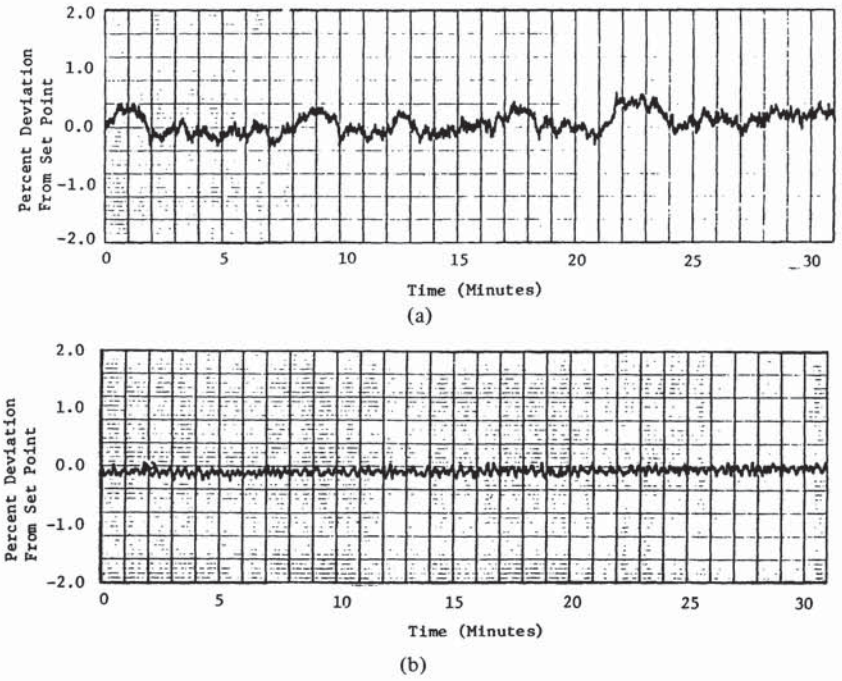

Fig. 4. (a) Steady-state power profile under analog control. (b) Steady-state power profile under digital control.

RESULTS OF THE EXPERIMENT

The digital control system has been tested under both steady-state and transient conditions that include natural phenomena such as xenon effects and operator-induced load changes. Sampling interval ranging from 0.2 to 0.5 s used for the tests. The following conclusions are deduced from the experimental results.

1) The system is tolerant of one or two failures in power, flow, and ΔT sensors. This fact has been verified by both natural and induced sensor failures that have occurred during the test runs.

2) The system is capable of calibrating the power sensors on-line during transients involving xenon variations, fuel depletion, temperature fluctuations, and operator-induced power-demand changes.

3) The controller maintains the reactor power at the desired level during both steady-state and rapid transient conditions, even though the power sensors themselves may be affected by process disturbances.

4) The digital controller is less sensitive to sensor degradation and noise than the original analog controller.

5) The system reliably exhibits on-line information on plant component and sensor failures and estimates of measured variables.

Under normal steady-state operations, data taken from one of the neutron flux detectors are shown in Fig. 4(a) and (b) to illustrate the relative performances of the analog and digital control systems. The digital controller maintains the reactor power steady within a much narrower amplitude band than does the analog controller. It appears from Fig. 4(a) that the analog controller causes the reactor power to oscillate at a frequency lower than that of the reactor noise. These oscillations can be attributed to the relay logic in the analog control law. The tiny ripples observed in the power profile under digital control in Fig. 4(b) are generated due to the inherent process disturbance. The continuous-action digital controller is capable of moving the regulating rod at a variable speed to compensate for part of the process noise within the passband of the controller filter. In contrast to the fixed gain, reset rate, and derivative time in the analog (P-I-D) controller, the digital controller has a variable gain which is automatically adjusted for the nonlinearities of the process gain, and the power unbalance (due to process disturbance) are dynamically compensated by the feedback of reactivity. Since the reactivity is calculated on-line via a nonlinear dynamic model of the physical process, it functions as a nonlinear compensator in the control law. The major reasons for the improvements under digital control can be attributed to the following: 1) the continuous-action control law replacing the relay logic; 2) a nonlinear compensator (involving gain adjustment and reactivity) instead of the fixed parameter linear (P-I-D) controller; and 3) sensor noise reduction in the computation of estimated power.

Fig. 5 shows the power output of the reactor during a test of the digital controller's response to power demand changes. (Similar tests were not conducted with the analog controller for operational safety.) The reactor's power output was initially 1.0 MW when the demand was increased to 4.0 MW in steps of 1.0 MW and then held stationary for about 22 min. The demand was then returned to 1.0 MW in similar steps by reversing the procedure. The reactor power closely followed the demand changes during each of the step disturbances shown in Fig. 5. During the load changes, the power sensors were automatically calibrated on-line to compensate for the distortions in the flux profile caused by changes in xenon distribution, fluctuations in coolant temperature, and nonlinearities over the power range. Practically, there were no

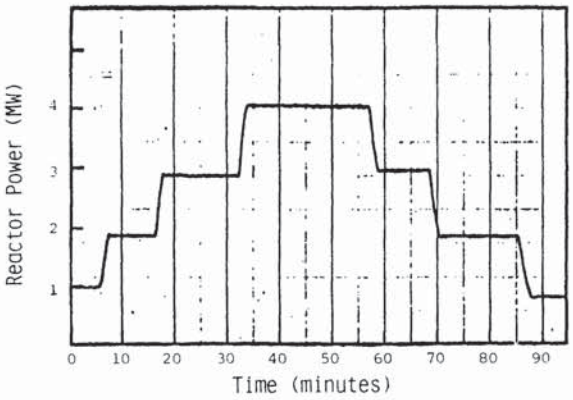

Fig. 5. Power transients for stepwise load changes under digital control.

overshoots or undershoots in the power profile during these operational transients. The single-input single-output controller described earlier is scheduled to be extended to a multivariable digital controller which, in addition to the reactor power, will regulate the average primary coolant temperature by automatic manipulation of the secondary coolant flow.

The fault-tolerant control scheme described here can be activated by one or more power sensors and is not restricted to specific control structures such as P-I-D in the analog controller. Functionally, this control scheme is more flexible and, from the point of view of sensor degradation and failures and the malfunctioning of other pertinent plant components (such as primary coolant pump and heat exchangers), more reliable and accurate than the original control scheme. Since the computer system in its present form consists of single-string components, i.e., there are no redundant CPU, memory, clock, A/D and D/A converters, the system hardware is not fault-tolerant. Therefore, the experiments have been designed such that possible hardware failures in the computer system should trip the controller to the manual mode. This is required by the MIT Reactor Safeguards Committee, and the details of the safety procedure are given in [16]. Replacement of the existing single-string hardware by a fault-tolerant computer system [20] such as FTMP [21] will lead to the evolution of an integrated control system which is tolerant of hardware and software failures of the computer and its ancillaries, sensors, and other plant equipment.

SUMMARY AND CONCLUSION

A fault-tolerant control system software has been developed and implemented using an LSI-11/23 microcomputer system. The control system has been demonstrated for on-line operations involving closed-loop control of power, fault diagnosis, sensor calibration, and information display in an NRC-licensed nuclear reactor. The control law uses heuristics and learning theory for emulating operator instructions and allows safe and reliable regulation of reactor power under steady-state and rapid load changes. The system software is structured for 1) implementation on commercially available microcomputers using either central or distributed processors, and 2) adaptation to complex industrial processes such as chemical and power plants for on-line fault diagnosis, sensor calibration, bumpless process control, and information display and monitoring.

Software reliability is enhanced by a systematic test procedure. Replacement of the existing single-string hardware by a fault-tolerant computer system will lead to the evolution of an integrated control system which is completely tolerant of failures of the computer and its ancillaries, sensors, and other plant equipment.

ACKNOWLEDGMENT

The author acknowledges the contributions of Dr. John A. Bernard, Dr. Mukund N. Desai, and Professor David D. Lanning in this research work. Special thanks to Mr. Paul J. Menadier who was responsible for the fabrication of electronic and mechanical hardware.

REFERENCES

[1] S. Osder, "Chronological overview of past avionic flight control system reliability in military and commercial operations," in *AGARDograph 224*, P. R. Kurzhals, Ed., AGARD-AG-224, 1977.

[2] J. J. Deyst and A. L. Hopkins, "Highly survivable integrated avionics," *Astronaut. Aeronaut.*, pp. 30–41, Sept. 1978.

[3] B. A. Barcley, T. G. Lenox, and C. J. Boxco, "Full authority digital electronic control—Highlights of next generation propulsion technology," *ASME Paper 78-GT-165*, Apr. 1978.

[4] G. M. Stanley, "On-line data reconciliation for process control," presented at the *Amer. Inst. Chem. Eng. Winter Annu. Meet.*, Nov. 1982. Paper 11b.

[5] D. I. Morris, "Computer control at Bruce nuclear generating station," *Nuclear Safety*, vol. 15, pp. 691–701, Nov.–Dec. 1974.

[6] J. L. Tylee, "A generalized likelihood ratio approach to detecting and identifying failures in pressurizer instrumentation," *Nuclear Technol.*, vol. 56, pp. 484–492, Mar. 1982.

[7] R. M. Clark and B. Campbell, "Instrument fault detection in a pressurized water reactor pressurizer," *Nuclear Technol.*, vol. 56, pp. 23–32, Jan. 1982.

[8] A. Ray, M. Desai, and J. Deyst, "Fault detection and isolation in a nuclear reactor," *J. Energy*, pp. 79–85, Jan./Feb. 1983.

[9] M. Desai and A. Ray, "A fault detection and isolation methodology—Theory and application," in *Proc. American Control Conf.*, 1984, pp. 262–270.

[10] A. Ray *et al.*, "Analytic redundancy for on-line fault diagnosis in a nuclear reactor," *J. Energy*, pp. 367–363, July/August 1983.

[11] C. H. Meijer *et al.*, "On-line power plant signal validation technique utilizing parity space representation and analytic redundancy," Elec. Power Res. Inst., Palo Alto, CA, Rep. NP-2110, Nov. 1981.

[12] J. L. Fisher *et al.*, "Signal validation for a BWR suppression pool," *Trans. Amer. Nuclear Soc.*, vol. 44, pp. 64–66, Aug. 1983.

[13] O. L. Deutsch *et al.*, "Development and testing of a real-time measurement validation program for sodium flowrate in the EBR-II,"

Charles Stark Draper Lab., Cambridge, MA, Rep. CSDL-R-1592, Oct. 1982.

[14] A. Ray and M. Desai, "A calibration and estimation filter for multiply redundant measurement systems," *J. Dynamic Syst., Measurement Contr., Trans. ASME,* pp. 149–156, June 1984.

[15] "Reactor systems manual," Mass. Inst. Technol., Cambridge, Rep. MITNRL-004, 1980.

[16] A. Ray, J. A. Bernard, and D. D. Lanning, "On-line signal validation and feedback control in a nuclear reactor," in *Proc. 5th Power Plant Dynamics, Control and Testing Symp.,* Mar. 1983, Paper 38.

[17] K. S. Navendra and M. A. Thathachar, "Learning automata—A survey," *IEEE Trans. Syst., Man, Cynbern.,* vol. SMC-4, pp. 323–334, July 1974.

[18] D. Waltz, "Artificial intelligence: An assessment of the state-of-the art and recommendation for future directions," *The AI Mag.,* pp. 55–67, Fall 1983.

[19] J. A. Bernard and A. Ray, "Experimental evaluation of digital control schemes for nuclear reactors," in *Proc. 22nd IEEE Conf. Decision and Control,* 1983, pp. 744–751.

[20] D. P. Siewiorek and R. S. Swarz, Eds., *Reliable System Design.* Bedford, MA: Digital Press, 1982.

[21] A. L. Hopkins, T. B. Smith, and J. H. Lala, "FTMP—A highly reliable fault-tolerant multiprocessor for aircraft," *Proc. IEEE,* vol. 66, pp. 1221–1239, Oct. 1978.

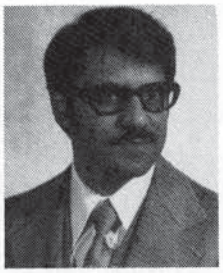

Asok Ray (SM'83) received the B.E. and M.E. degrees in electrical engineering from Calcutta University, and the M.S. degree in mathematics and the Ph.D. degree in mechanical engineering from Northeastern University, Boston, MA.

He served on the faculty of Mechanical Engineering at Carnegie-Mellon University and has held research and management positions in industry. Currently, he is an Associate Professor of Mechanical Engineering at Pennsylvania State University, University Park. He is also affiliated with the Massachusetts Institute of Technology, Cambridge. His research interests are in the mainstream of control and instrumentation, and modeling and simulation of dynamical systems with diverse systems engineering applications. He is the author or coauthor of more than 60 research publications in archive journals and reputed conferences.

Dr. Ray is a Registered Professional Engineer in the Commonwealth of Massachusetts. He is also a member of ASME and of Sigma Xi.