CONF-910535--2

CONF-910535--2

DE91 009637

Instrumentation and Controls Division

PLANT MONITORING AND SIGNAL VALIDATION AT HFIR*

James A. Mullens
Oak Ridge National Laboratory, Oak Ridge, TN 37831-6010, USA

Paper to be presented at the
Sixth Specialists Meeting on Reactor Noise
Gatlinburg, Tennessee, USA
May 19–24, 1991

*Research sponsored by the U.S. Department of Energy under contract DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc.

PLANT MONITORING AND SIGNAL VALIDATION AT HFIR*

James A. Mullens
Oak Ridge National Laboratory
Instrumentation and Controls Division
Oak Ridge, Tennessee 37831-6010

ABSTRACT

A monitoring system has been developed for a nuclear reactor at Oak Ridge National Laboratory. The system provides real-time monitoring at the plant and at remote sites. The system is used also to test reactor monitoring ideas including signal validation schemes and operator aids.

INTRODUCTION

This paper describes a monitoring system for Oak Ridge National Laboratory's (ORNL's) High Flux Isotope Reactor (HFIR). HFIR is an 85-MW pressurized water reactor designed to produce isotopes and intense neutron beams. The monitoring system performs the following functions:

(1) logs plant signals to provide a permanent record of plant operation,
(2) displays plant status in the shift supervisor's office and at remote sites,
(3) alerts its users to abnormal plant conditions, and
(4) replays logged HFIR operation for analysis.

PLANT SIGNALS AND COMPUTER SYSTEM

The plant supplies 132 discrete (on/off) signals and 49 analog signals to the monitoring system through isolation devices. Most signals come from the primary system, the coolant cleanup systems, and radiation monitoring.

The system uses a CAMAC [1] data acquisition system and a VAX computer. The CAMAC analog signal acquisition card can provide measurements every 2 s and remove 60-Hz contamination. The computer is a micro-VAX-II/GPX with 16 Megabytes (MB) of RAM, 216 MB of disk storage, a mouse, a keyboard, and a GPX high-resolution (1024 × 864), 256-color graphics display. The display is in the HFIR shift supervisor's office, adjacent to the control room. The system is linked to the ORNL computer network via a dedicated 56-kilobaud phone circuit connection to a VAX in the Instrumentation and Controls Division surveillance and diagnostics laboratory.

The computer runs under VMS, a multiuser operating system which, while not a full-feature, real-time operating system, does provide task priorities so that time-critical tasks can be performed ahead of other tasks. VAX Workstation Software (VWS) windowing software and an implementation of the GKS [2,3] graphics library drive the GPX color display. DECnet software performs all network communications.

*Research sponsored by the U.S. Department of Energy under contract DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc.

MONITORING OVERVIEW

The plant computer measures, logs, and distributes plant signals over the computer network, and runs the full monitoring system. Remote computers also run the full monitoring system.

The monitoring system first does some signal checking and sets signal validity flags which are considered whenever the signal is used in a calculation or shown to the operator. Some generally useful calculations are then performed (e.g., conversion to engineering units, averaging redundant signals, time-averaging signals, and computing noise on the signals). A series of modules then examines the new data. These modules could be loosely classified as additional signal validation, plant configuration checks (e.g., pump and valve lineups), operational boundary checking (e.g., a temperature approaching an alarm limit), and operator aids for routine operations (e.g., marking which parts of the plant are undergoing changes). Finally, the graphical status display is updated and commands from the operator are processed.

DATA ACQUISITION, LOGGING, AND NETWORK DISTRIBUTION

The signal acquisition and logging system (SALS) provides analog data every 2 s and polls the discrete signals (i.e., alarms and relays) every 1 s. The system was not intended to perform signal analysis, so no investment in equipment to perform Nyquist filtering was made. SALS writes every analog and discrete measurement to a shared memory buffer which holds all measurements made in the previous hour, with their timestamps. The shared memory buffer is available to all programs needing plant measurements.

SALS also logs the measurements to files. SALS writes analog data periodically and discrete data as they change. The analog logging rate can vary within a log file, and SALS will change the logging rate on a command from the operator or other software running on monitoring system.

The DECnet computer network provides access to log files from remote computers and a means for programs running on remote computers to communicate with the monitoring programs running on the plant computer. This latter capability provides distribution of current HFIR signals from the plant VAX/CAMAC system to remote machines. This capability allows (1) remote emergency monitoring of the plant and (2) development and testing of new monitoring software before installation at the plant. Users monitor the HFIR remotely by running the full monitoring system on a remote VAX.

SIGNAL VALIDATION

The author has two major goals for signal validation at HFIR. First, there must be a very low false alarm rate, especially during unusual plant conditions so that these alarms do not add to the confusion of a plant transient. This requirement requires careful design and testing. Second, validation would ideally be done the same way an operator would do it so that validation alarm messages can be readily interpreted by the operators and justified to plant management. It would be less desirable to implement a validation scheme that was so abstract that plant operators could not understand how it works. The general strategy is to validate by using simple relationships between the signals, detect only faults severe enough to be of concern to the operator, and take action only if discovered faults are persistent.

Sensor Information

All analog signals have a nominal range of 1- to 5-V signals. Any signal outside this range is immediately flagged as bad. Similarly, some sets of discrete signals can be self-contradictory. For example, a valve position monitored through limit switches cannot be fully open and fully closed at the same time.

The most crucial measurements (i.e., primary flow, vessel inlet temperature, vessel outlet temperature, flux, and reactor period) are made with three redundant sensors. These sensors can be compared against one another to detect failures in all conditions except fast plant transients (which can give false alarms because of slightly different dynamic responses and sensor locations). The three-signal case will be described in detail before a more general scheme is presented.

The difference between each pair of signals is computed (3 differences) and judged to be high or satisfactory by comparing it with the instrument calibration specification. The eight possible results of these three comparisons lead to four distinct situations. First, all comparisons are satisfactory: nothing more needs to be done. Second, none of the signals compare well: no evidence exists to invalidate any particular signal, but at least two must be wrong so they are all invalidated. Third, one signal disagrees with the other two: the errant signal is invalidated. Fourth, a signal with a median value compares well with the signal having a high value and the signal having a low value, but the high and low signals disagree with one another: this situation is ambiguous. Validating the middle value is adequately justified because its value compares well with the other two signals, and any reasonable method of estimating the true value of the measured quantity will tend to give the middle value. The problem is what to report to the operator. The chosen solution is to mark invalid the signal that has the largest total disagreement with the other signals. If the problem persists and in subsequent plant measurements this signal continues to be the worst, then the initial diagnosis was probably correct. If subsequent measurements cause the other signal to be marked invalid instead, this is a clear indication to the operator of the ambiguity of the situation.

This scheme can be implemented for an arbitrary number of signals by iteratively invalidating the worst signal and dropping its comparisons until a consistent set of signals remains (or no signals remain). It can be used also to compare computed estimates of a signal's value based on other signals (e.g., using reactor vessel pressure drop to estimate primary flow and comparing that to the primary flow sensor). This scheme resembles the parity space technique [4], but the application is different because the level of failure being detected is severe enough that any bad comparison is a reportable failure.

As an example of an application, the period signals are so noisy that the period estimates during plant startup are difficult to interpret. A better estimate results from "failing" a signal according to the above criteria and using the other two signals. Figure 1 illustrates the improvement in the period estimate.

The monitoring system is not given warning when instruments are taken off-line for tests. The redundant signal validation allows the monitor to handle this problem automatically so that the operator is not being shown obviously wrong information and the monitoring functions are not being fooled into issuing alerts to the operators.

Plant Models

Many opportunities are available to validate signals by comparing them through simple models of their behavior. Experiments are being carried out for several such models, with initial emphasis on reporting discrepancies and signals involved to the operators. These models currently build on the reliable validation provided by the redundant sensors to extend validation to single-sensor measurements. For example, a sustained change of the reactor vessel pressure drop without a corresponding primary flow decrease indicates a failed vessel pressure drop signal.

One modeling effort involved the redundant sensors. Plant measurements were used to develop a model of the calibration difference between the signals and the typical disagreement remaining after accounting for calibration. The monitor then performed redundant signal validation on the modeled signals (using the modeling errors as limits) instead of the raw measurements (using the plant calibration specifications as limits). The advantage, which was proven with plant data, is more sensitivity to changes in instrument behavior. The disadvantage is that normal calibration drifts during a fuel cycle were large enough that signals were soon being reported as failed. An adaptive calibration was considered; however, the plant normally operates in a narrow range of conditions so that calibration could be adapted on only a narrow range of the instrument's scale. As soon as a serious transient occurred and signals moved outside this range, reports of signal failures would only add to the confusion and cause a loss of confidence in the system. Plant operators

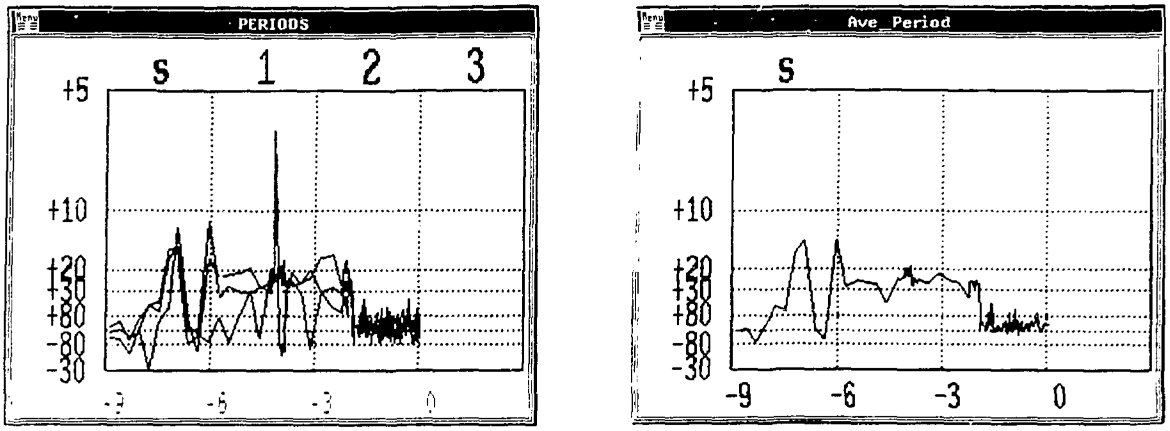

(a) *(b)*

Fig. 1. Reactor period signals (in seconds) during initial criticality and startup, showing *(a)* the three measured signals and *(b)* the result of removing invalidated signals and averaging.

are not interested in on-line reporting of normal calibration drifts, but this scheme is being considered as a tool for instrument technicians.

An on-line application examines the state of the primary coolant system. It compares the primary flow rate to the flow rate that should result from the pump and valve lineups reported by plant signals. A discrepancy, an unusual flow value, or an unusual pump/valve lineup will cause the monitor to execute a rather complicated procedure based entirely on plant-specific knowledge. It considers several factors such as how closely the flow rate matches the typical flow rate, whether the flow rate correlates with the pressure drop across the reactor vessel, the disagreement among the three primary flow sensors, valve limit switches report a valve simultaneously open and closed, a pump reported on in an isolated flow loop or off in an open loop, and a heat exchanger whose secondary valve lineup does not match its primary valve lineup. The procedure estimates the most probable choice of (1) measured flow is correct but unusual (plant equipment failure indicated), (2) a flow rate signal has failed, (3) pump/valve lineup is unusual but is being reported correctly, or (4) a pump/valve signal has failed.

Because the plant operators do not monitor the plant signals' noise behavior, this is one area in which the monitor can make a unique contribution. The most convenient way to estimate signal noise in this monitoring system is to use exponential averaging. The average value of a signal at sample time t is given by:

$$A(t) = \{1-f\}\ S(t) + f\,A(t-1)\ ,$$

where A = exponential average, S = signal value, and f = exponential memory (adjusted to the varying sampling interval to yield a consistent averaging time constant). The noise estimate used is

$$N(t) = \{1-f\}\ |S(t) - A(t-1)| + f\,N(t-1)\ ,$$

where N = noise estimate. The noise is based on the absolute value (rather than the square) of the difference between the current signal value and the average signal value. Squaring strongly emphasizes outliers and makes $N(t)$ for some signals very unstable. Noise levels are calculated for all measured signals during monitoring.

The reactor period signals are examples of the use of this noise. Their fluctuations are largely random and not strongly correlated. The model (plant information) for signal validation consists of (1) a minimum noise level that should exist, (2) a maximum noise that should exist, (3) and specifications of the conditions under which should be no reactor period or period fluctuations due to refueling activities, and (4) a maximum allowable disagreement between the period signals. In March 1990, a broken wire at HFIR resulted in the loss of the flux signal to a reactor period channel. Figure 2 shows the period signals that resulted and the detection of this problem by monitoring for a minimum noise level in the signal. The monitoring system generates an alarm at time 0 on the chart.

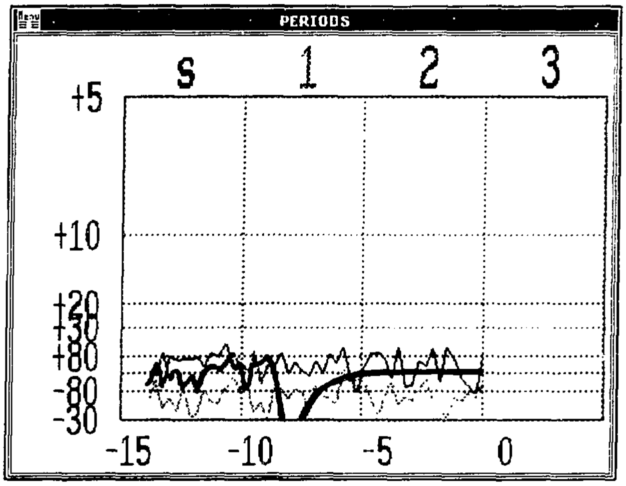

Fig. 2. Reactor periods during shutdown conditions, measured every 16 s, showing failure of one channel (thick line). Monitor alarmed at time 0 on the chart.

Minimum noise level monitoring is also active for signals derived from pressure measurements. The anticipated failure is a blockage of the sensing line which would lead to a quiet signal that is frozen but may appear correct.

<u>Use of Validation Results</u>

The monitoring system has a signal quality flag for every plant signal, and these quality flags are propagated to quantities calculated from the signals. For example, there are three channels of reactor vessel inlet temperature, outlet temperature, and primary flow. Each channel is used to compute a thermal power value. If channel 1 inlet temperature is invalid, its failure flag is propagated to the channel 1 thermal power estimate. When the average thermal power is calculated from the channels, only channels 2 and 3 will be used and the result will be flagged as "good" because good data were available.

STATUS DISPLAYS

The monitoring system provides two types of displays, a graphical display and a computer terminal display (text only). The graphical display requires a VAXStation, but the terminal display allows only an ANSI-standard terminal.

Graphical Interface Overview

The plant status display is divided into several windows: a large "plant schematic" window, a message logging window, and up to ten small stripchart windows. Every window can be moved, resized, or shrunk to an icon.

The schematic window is the main display; it contains a diagram that illustrates some subsystem of the plant. It currently shows one of four views: the primary system, the secondary system, the primary coolant cleanup system, or the plant radiation alarms. The primary system view (Fig. 3) also displays most signals received, allowing the operators to use it in most circumstances as an overview of the plant.

The strip-chart windows emulate strip-chart recorders and can show several signals simultaneously. These windows pop up when the operator requests them or when the monitoring system has detected that the signal's value is changing more than usual.

The operator interacts with the status display by using only the mouse. A mouse click on an object in a graphical window causes the system to display a menu related to the object. For example, clicking on a temperature reading brings up a menu allowing the operator to launch a strip chart of the temperature, view information describing how the temperature has been changing, etc. The operator then clicks on the menu to select anything of interest. The system continues updating the plant status display in the background while the operator is selecting items from menus in the foreground.

Only 8 colors are currently used in the display, although 256 are available. More are not necessary because people cannot quickly distinguish many more than eight colors [1, p. 338; 2, p. 5-10]. Additional colors would produce a more attractive interface but may not give more information to the operator. The colors are used consistently for the following purposes:

- red — alert highlighting: operator's attention needed;
- yellow — warning: alert condition acknowledged by operator but persistent;
- green — operating or ready: component operating or in its normal state;
- black — not operating or not ready: component shut down or not in its normal operating state;
- blue — functional but no operational states defined for this component;
- magenta — radiation alarm on;
- cyan — miscellaneous items such as outlines on component drawings; and
- grey — window background and text labels on components.

The display software is written in the C++ language [3], which is a strongly typed, object-oriented language well suited to handling graphical user interfaces. Each object appearing in the window is one of the classes of objects shown in the family tree depicted in Fig. 4. Each window is under the control of a display manager (another type of object) that manages the collection of objects displayed on the window and processes operator input.

Display of Signal Values and Alert Status

Each object associated with plant signals displays its signal value and its alert state. The signal value display varies with the component. For example, heat exchangers are green when they are ready to exchange heat and black otherwise; a digital meter displays a signal value with numerals; and a strip chart plots the history of a signal. The alert state is shown by bold red or yellow highlights if the state is alert or warning. No highlights appear if the state is normal.

Most objects display a quantity derived from plant measurements rather than the raw plant measurement itself. For example, the displayed reactor vessel inlet temperature is the average of the three channel measurements; and heat exchangers change color to show whether the heat exchanger is ready to operate, which depends on four valves with eight plant signals ultimately used to make a determination.

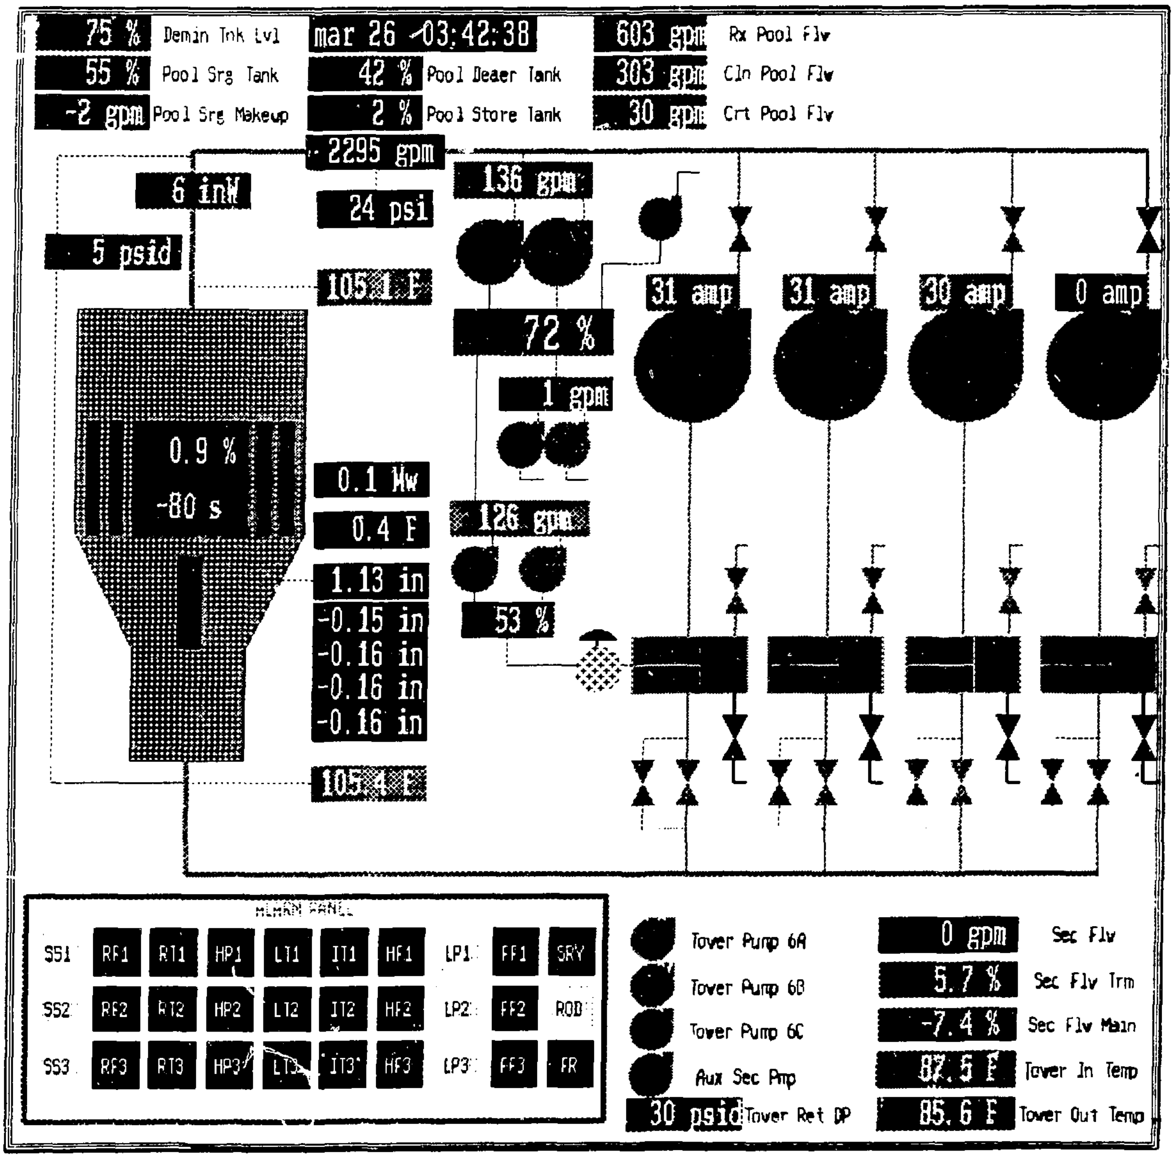

Fig. 3. The main view of the status display concentrates on the primary system but shows major components of other systems as well (original in color).

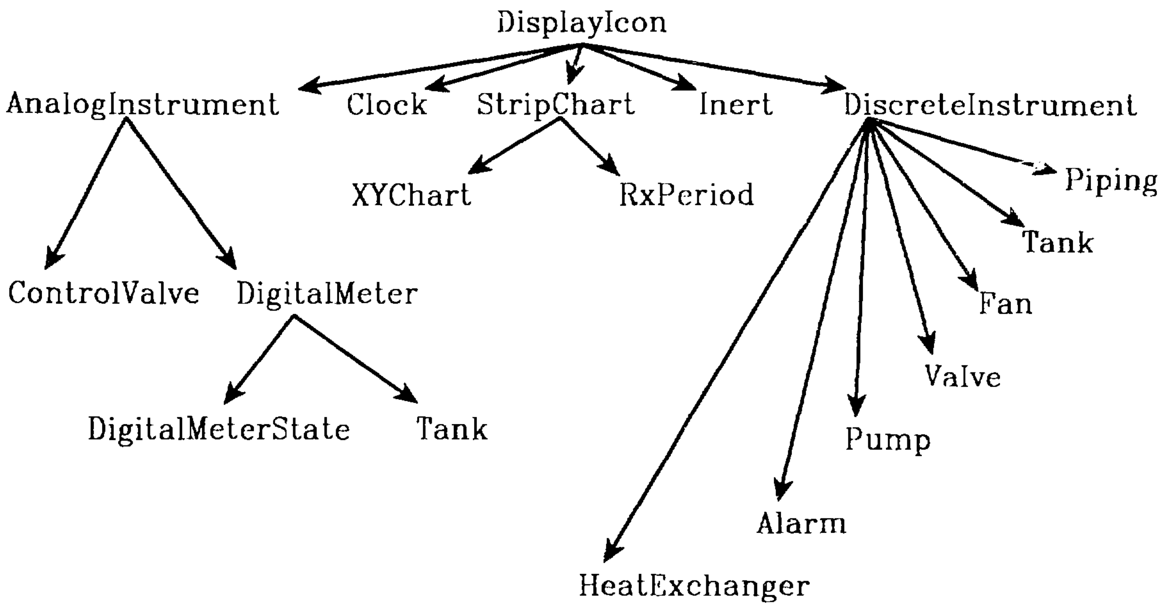

Fig. 4. The family of objects that appear in the graphical interface.

If signals are declared bad by the signal validation algorithms, display objects replace the value display with a bad-signal display. For example, the numerals of the DigitalMeter object are replaced by the characters "*BAD*." However, the operator can see the bad value by interrogating the object. In this way, misleading information is not shown but the operator can still investigate signal problems.

Interaction with the Operator

When the operator clicks on an object, the operator begins a dialogue with the object. The object always responds with a menu containing a variety of information and actions. Figure 5 shows the menu for a DigitalMeter object (typical for objects representing analog signals). Action items are written with the first letter of each word capitalized, and they are phrased in the command voice. Actions items will display more detailed information or alter operation of the monitoring system. Information items describing abnormal conditions are written in uppercase letters (e.g., "status: ALERT"), and items showing normal conditions are written in lowercase letters.

Strip Charts

Strip charts can be created in one of two ways. The operator can request a chart, or the monitoring system will create a chart automatically when a signal has changed significantly.

A typical strip chart is shown in Fig. 6. The strip charts are a valuable part of the display because they provide a way to investigate plant transients. The operator can create a chart at any time for any displayed analog signal or quantity the monitor calculates and for any duration between 2 and 60 min. A chart, when created, shows the recent history of the signal and updates continuously. The charts stack themselves vertically so that their time axes are aligned to aid transient analysis. They can show multiple signals (on one common y-axis scale).

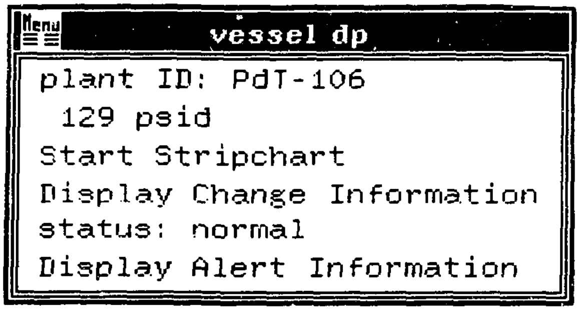

Fig. 5. Selecting an object in a display window always causes a menu to appear listing information and actions the operator can perform.

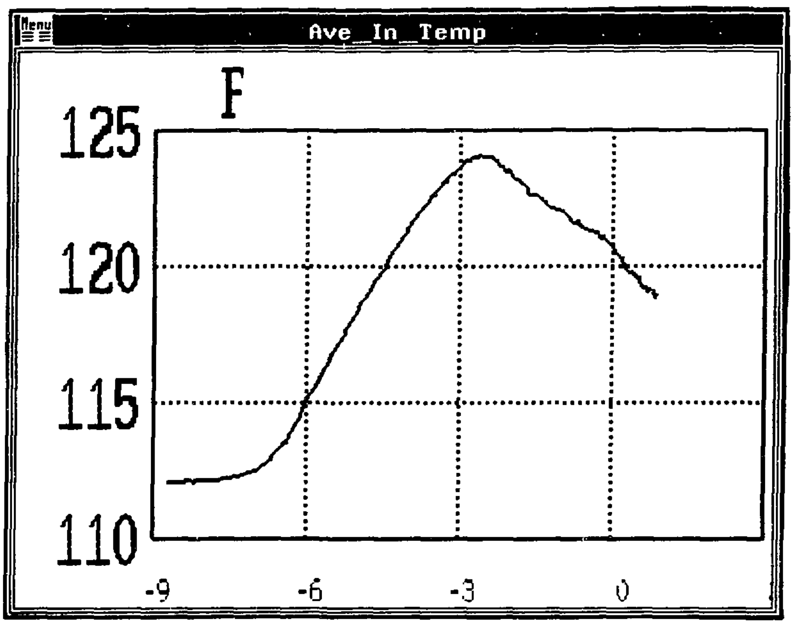

Fig. 6. Strip chart windows scale automatically, initially show recent history (up to 1 h), and label time in minutes (2 to 60).

A strip chart dynamically scales itself to changing signals. The range of the chart of a single signal is the larger of (1) 2% of the signal range, (2) the range the plant-change monitoring system has established as the normal range of fluctuation of the signal, or (3) the actual variation of the signal during the time displayed. When a chart shows multiple signals from a redundant sensor group (e.g., flux) the chart will scale as tightly as possible on the assumption that the operator is interested in the small differences between the signals.

If interested in a point on the strip chart, the operator can click on the point to see the time (hours:minutes:seconds) and the values of the signals at that time.

CONDITION MONITORING

Limiting Safety System Settings

The limiting safety system settings (LSSS) that the monitor checks are a group of temperature and pressure set points at which action should be taken to avoid unsafe conditions. Fourteen primary pressure set points exist for such actions as rupture disk opening and primary pump shutdown. Nine primary coolant temperature set points are available for actions such as high-temperature scram and forced depressurization at low temperature. These set points are monitored, and problems are logged and alarmed with a warning bell.

Cooling System Operation

The monitor tracks the valve and pump lineups for primary and secondary flow, the rate of change of primary temperature, and the position of the secondary flow control trim valve (which is under automatic control) and issues warnings as appropriate. Specifically, the monitor watches for (1) main pump on with pump isolation valve shut, (2) no heat exchangers with flow on both the primary side and the secondary side (bad valve lineup), (3) reactor vessel inlet coolant temperature projected to exceed limits within 60 s, and (4) secondary flow control trim valve projected to exceed its effective control range of 80% open or 80% closed within 60 s.

OPERATOR AIDS

Secondary Flow Control

The operator must manually adjust the position of the main secondary flow control valve during power maneuvers. This task is complicated by the slow response of the plant to the valve changes and the operation of an automatically controlled secondary flow control trim valve. Purely as an experiment, the monitoring system checks the balance between primary and secondary cooling and issues an alert if it judges that the control valve should be adjusted. This application is complicated by a lack of good signals from the plant secondary.

The monitoring system does not receive the heat exchanger secondary coolant temperatures, and the secondary flow signal is not accurate at the low flow rates that occur during plant startup. This means that the monitoring system cannot directly measure secondary cooling or estimate the proper secondary flow control valve setting. To judge the situation, the monitoring system uses the position of the manually adjusted main flow control valve, the position of the automatically controlled trim flow control valve, the primary coolant temperature, the rate of change of primary coolant temperature (which indicates an imbalance between primary power and secondary cooling), and the speed at which valves and temperatures are approaching the limits of their allowable operating ranges. The values of these signals are classified as "high", "low", etc.; and these classifications are used in a series of rules (boolean logic) to determine whether the valve should be repositioned and the urgency of doing so. This information is diagnostic only because the monitor is not allowed to affect the operation of the plant. The next step in this experiment is to add a "schematic" view to the plant status display dedicated to the handling of secondary flow.

Plant Change Monitoring

CHGMON (CHanGe MONitor) alerts the operator to changes in plant conditions. It is most useful during steady-state plant operations and is "tuned" to monitor full-power, steady-state operation.

When CHGMON starts, it takes the current plant condition as the "reference" condition and begins monitoring for changes. The plant condition is specified by the current values of the displayed plant signals, averaged for 10 consecutive measurements. CHGMON assumes that ideal plant behavior is small fluctuations of these signals around their reference value.

CHGMON set bounds around the signal value for the detection of trends and large changes in the signal value. When the signal value crosses the trend boundary, a trend counter is started. If the signal remains outside the trend boundary until the counter has counted to its limit, a trend notice is issued. A value change alert occurs as soon as the signal crosses the wider "value change" boundary. Either change alert will produce a strip chart to illustrate the change.

CHGMON has a tracking mode in which it will alert the operator to a change in a signal value and then reset the reference value of the signal to the new value. This action allows CHGMON to adjust itself to changes in plant operation while informing the operator of the progress of these changes. For example, draining a tank will cause of series of "alerts" as the tank level drops.

Figure 7 shows a menu that gives information about changes in any displayed signal. The first item shows the status (either "unchanged" or "ALERTED"), and the second and third items show the alerted signal value and time of the alert. Other menu items summarize the change-monitoring limits and the behavior of the signal. The first and last items of this group are the change limits currently in effect (i.e., the highest and lowest signal values *allowed*). Adjacent to the limits are the highest and lowest values *observed*. If the signal has exceeded a limit, the menu item is capitalized. The current value of the signal is in the middle of this group. Other items summarize the trend monitoring including the trend limit, the current trend count, and the longest trend observed. The last items allow the operator to tune change monitoring for a signal. "Reset Nominal Value" resets the reference signal value to the current value, and "Adjust Change Allowed" causes the monitor to accept the previously observed fluctuations of the signal as normal and reset the current signal value and change limits accordingly.

Another program, SIGSTATS, estimates limits for the change detection done by CHGMON. It reads a plant data log file and calculates a variety of statistics describing plant signals. Its main purpose is to separate nonstationarity of the signal mean from fluctuations in the signal. Stationarity is usually tested at 15-min intervals, mainly because of an oscillation between primary and secondary heat balances with a period of up to 5 min. The resulting change limits represent normal plant behavior for the plant condition that existed in the log file with the plant operating in steady state. These limits represent statistically significant changes (i.e., a signal will rarely exceed them if the plant is in steady state), but these are small changes from the operator's point of view. A coarser set of limits is provided for the plant monitor.

CHGMON's value during major plant upsets is that it automatically starts five-minute strip charts on the signals first observed to change. Its value during normal operation is that it marks activities throughout the plant on a single display. The pattern of red change highlighting on a plant schematic display can be understood in a glance to mean activity in a particular subsystem of the plant.

CONCLUSIONS AND FUTURE WORK

A monitoring system for ORNL's HFIR reactor has been developed and is being used at the plant and at remote sites. Signal validation and diagnostic functions are being added as they are proven to work with a minimum of false alarms. This process requires careful design that considers many limitations: pertinent plant signals not accessible, imperfect plant instrumentation, data acquisition system capabilities, low computer performance, operator preferences, and the need to avoid false alarms during unusual plant situations.

Future work will include the addition of more plant signals, more signal validation and diagnostic capabilities, conversion of the status display to a windowing system based on the X standard, integration of the system with the RELAP plant simulation software and its graphical interface, implementation of displays designed by plant operators, and a change monitoring display option to suppress highlighting of anticipated changes so that only initiating or unusual events are highlighted on the display (alarm filtering).

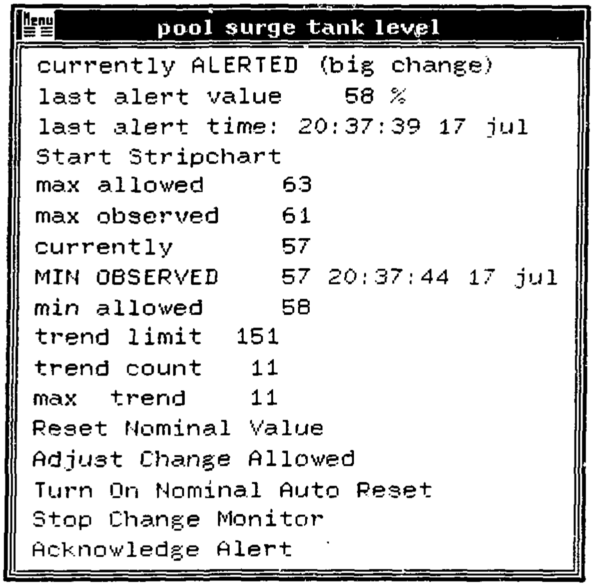

Fig. 7. The signal change monitoring pop-up menu shows a variety of information about the signal's behavior and allows the operator to refine change monitoring for the signal.

REFERENCES

1. *Modular Instrumentation and Digital Interface System (CAMAC)*, ANSI/IEEE Std. 583-1982, American National Standards Institute, New York, 1982.

2. *ISO/DIS 7942 Information Processing—Graphical Kernel System (GKS)—Functional Description: GKS Version 7.2*, ISO/TC97/SC5/WG2 N163, 1982.

3. F. R. A. Hopgood et. al., *Introduction to the Graphical Kernel System (GKS)*, Academic Press, New York, 1983.

4. "Thresholdless Redundancy Management with Arrays of Skewed Instruments", *Integrity in Electronic Flight Control Systems*, AGARDograph #224, 1977.

5. B. Shneiderman, *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, Addison-Wesley Publishing Company, Reading, Mass., 1987.

6. P. R. Frey and W. H. Sides, *Computer-Generated Display System Guidelines*, NP-3710, Electric Power Research Institute, Palo Alto, Calif., 1984.

7. B. Stroustrup, *The C++ Programming Language*, Addison-Wesley Publishing Company, Reading, Mass., 1986.

8. D. M. Snider, K. L. Wagner, and W. Grush, *Nuclear Plant Analyzer (NPA) Reference Manual*, EGG-EAST-9096, EG&G Idaho, Inc., Idaho Falls, 1990.