Reconfigurable Control of Power Plants Using Learning Automata

Humberto E. Garcia, Asok Ray and Robert M. Edwards

A deaerating feedwater heater, equipped with a water level controller and a pressure controller, has been chosen to investigate the feasibility of a reconfigurable control scheme for power plants by incorporating the concept of learning automata. Simulation results based on a model of the Experimental Breeder Reactor (EBR-II) at the Argonne National Laboratory site in Idaho are presented to demonstrate the efficacy of the reconfigurable control scheme.

Introduction

Steam electric power plants require integration of a large number of subsystems and controller modules, and the wide range of operating conditions lead to complex maintenance problems [1]. In addition, plant equipment failures and malfunctions must be rectified following prescribed safety and emergency procedures. Since human operators may not be able to take corrective actions for all possible malfunctions in a timely manner, it is desirable to have automatic responses to unanticipated events under normal operating conditions as well as during plant start-up and shutdown. One approach is to identify a malfunction on-line and then perform necessary engineering analyses to decide the best action. However, decisions based on routine applications of engineering analyses may not produce results on a timely basis. The situation becomes more complex with unstructured disturbances. An alternative approach is to apply the concept of learning automata [2],[3] to continually monitor the system performance. In this approach, the control structure can be reconfigured, and the individual learning modules are made to be parts of the integrated control and decision-making system which dynamically responds to changes in the plant conditions. This time-dependent response results from the enforcement of an appropriate control law selected from a bank of pre-designed controllers [4],[5]. The possible control laws could be formulated using the time-domain and frequency-domain analytical techniques [6],[7], fuzzy logic [8] or simple heuristics. The key idea is that no single controller can satisfy the performance requirements for all conditions of plant operations. The learning agent evaluates the system performance and reconfigures the system by switching to the controller that it determines to be most likely to achieve the desired performance at the given instant. However, poor performance may result due to incorrect identification of the plant operating conditions, imprecise knowledge of the plant, or inadequate control capability. In concept, the reconfigurable control system should possess the capability to deal with these uncertainties and function in a stochastic setting.

A learning automaton interacts with its environment and enables the system to update actions based on environmental changes. The learning procedure can be divided into two sequential steps. First, the automaton chooses a specific action from the finite set of actions offered by the environment. Second, the automaton is penalized or rewarded depending on the environment response, and thus influences the choice of future action. In the present context, a specific action corresponds to one of the several controllers, and the automaton learns which controller is most appropriate for a given plant condition.

The principle of stochastic automaton operating in an unknown environment has been briefly described in detail in Narendra and Thatachar [2],[3]. The automaton and the environment interact with each other in the same way as a plant and its feedback controller do. That is, at every sample, the action of the automaton is an input to the environment which generates a response to excite the automaton. A variable structure stochastic automaton (VSSA), used in the proposed reconfigurable controller, essentially updates the action probability vector at every sample such that the characteristics of the environment are specified by a set of penalty probabilities having a one-to-one correspondence with the set of actions generated by the automaton. On the basis of the environment response and the internal states of the automaton at each sample, the objective is to select the action which yields the minimum penalty.

System Description

A deaerating feedwater heater, equipped with a water level controller and a pressure controller, has been chosen to investigate the feasibility of a reconfigurable control system based on learning automata to deal with power plant operations. Besides removing entrained air from the feedwater supply system, a deaerating heater provides an inventory of water for the main feedwater pump(s). A power plant deaerator is elevated relative to the feedwater pump inlet to provide a net positive suction head [9]. Due to this physical separation, the effects of changes in pressure and temperature in the deaerator water do not appear at the pump immediately but are separated by a transport delay. Specifically, pressure waves travel at the speed of sound through water while changes in water temperature are transported at a speed proportional to the feedwater flow rate. The interconnecting pipe is short enough to neglect pressure delays but long enough for the temperature delays to be significant for plant design and operation. These differences in transmission of events may diminish the available net positive suction head during transients. A severe loss of the suction head could

---

Presented at the ASME Winter Annual Meeting at Dallas, TX, in November 1990. Humberto E. Garcia is with Electrical Engineering Department, Asok Ray is with Mechanical Engineering Department, and Robert M. Edwards is with Nuclear Engineering Department of the Pennsylvania State University, University Park, PA 16802. This work was supported by the U.S. Department of Energy (DoE) under Cooperative Agreement DE-FG07-89ER12889, Modification No. A000. The findings, conclusions, and ***recommendations expressed are those of the authors and do not necessarily reflect the views of DoE.***

0272-1708/91/0100-0085 $01.00 © 1991 IEEE

ruin the feedwater pump due to cavitation possibly within a few minutes [9].

To avoid any such potential damage of the feedwater pump, both deaerator pressure and its decay rate must be controlled. The deaerator pressure is normally maintained by the flow of high quality steam bled from the main turbine into the deaerator while the water level is regulated by adjusting the flow of condensate water from the low pressure feedwater train. If the steam inlet flow is abruptly reduced, the deaerator pressure would rapidly decay which may cause cavitation at the feedwater pump. A rapid pressure decay can be arrested by reducing the flow of the relatively cool condensate water into the deaerator. At the same time, the water level in the deaerator must not drop below a minimum limit to ensure that the feedwater pump does not run dry. During malfunctions the reconfigurable control system must decide which action should be taken to manipulate the condensate water valve, i.e., whether to adjust the condensate flow to maintain the water level in the deaerator tank or to prevent the pressure decay as much as possible.

The deaerator system of the Experimental Breeder Reactor (EBR-II) at the Argonne National Laboratory site in Idaho has been chosen to demonstrate (via simulation experiments) the efficacy of reconfigurable control based on learning automata. The deaerator is a vertical direct-contact tray-type heater with an integral storage tank having a capacity of 3200 gallons at the normal operating water level, and provides a feedwater capacity of 265 000 lbm/h at 150 psig and saturation conditions (364°F) at the deaerator outlet. During its normal full-power operation, the deaerator receives condensate flow of 188 000 lbm/h at 255°F, steam from the main steam header at a rate of 22 000 lbm/h and 727°F, and drain water flow from the high pressure heater. The deaerator is elevated to provide a net positive suctiion head of 12 psig at the inlet to the feedwater pumps.

The Reconfigurable Control

This section introduces the concept of a reconfigurable control scheme that is capable of adjusting itself to changes in plant operating conditions by incorporating learning automata [2],[3]. The functions of the reconfigurable control system include: i) identification of the plant operating condition; ii) evaluation of the current control performance; and iii) updating of the individual controllers' performance, followed by selection of a specific control law from the set of available ones. Several indices are combined to evaluate the performance of the selected controller. The controller selection process is based on the knowledge acquired by the system from the current and previous plant responses. A certain degree of intelligence is required to integrate these tasks. To this effect, learning behavior is incorporated in the reconfiguration scheme, and stochastic automata are specifically chosen because the plant is likely to undergo operational uncertainties.

Fig. 1 shows the reconfigurable control system where i) the deaerator is the plant; ii) condensate water flow is the control input manipulated by a continuous-position valve; and iii) exactly one of the two plant outputs, deaerator water level or pressure, are controlled at any given instant of time, i.e., the level and pressure controllers are not allowed to act simultaneously on the condensate flow valve. (Note: The reconfiguration scheme can be expanded to more than two controllers.) Both controllers are designed using the single-input single-output proportional-integral algorithm, and they always track each other such that bumpless switching from one controller to another is assured. The proportional gain and reset time of both controllers were determined on the basis of a simplified model of the deaerator, and are listed in Table I. Although more advanced algorithms [6]-[8] are likely to improve the system dynamic performance, these simple controllers facilitate a clear understanding of the operations of the learning mechanism.

The master module in Fig. 1 is the decisionmaking component of the reconfigurable control system, and makes use of other plant variables, such as feedwater flow, feedwater and condensate water temperatures. The steam valve which admits steam into the deaerator and is normally used to regulate

Table I: Controller Parameter Settings

Parameter Description	Controller #1 (Level)	Controller #2 (Pressure)
Proportional Gain	48	40
Reset Time	15 s	15 s

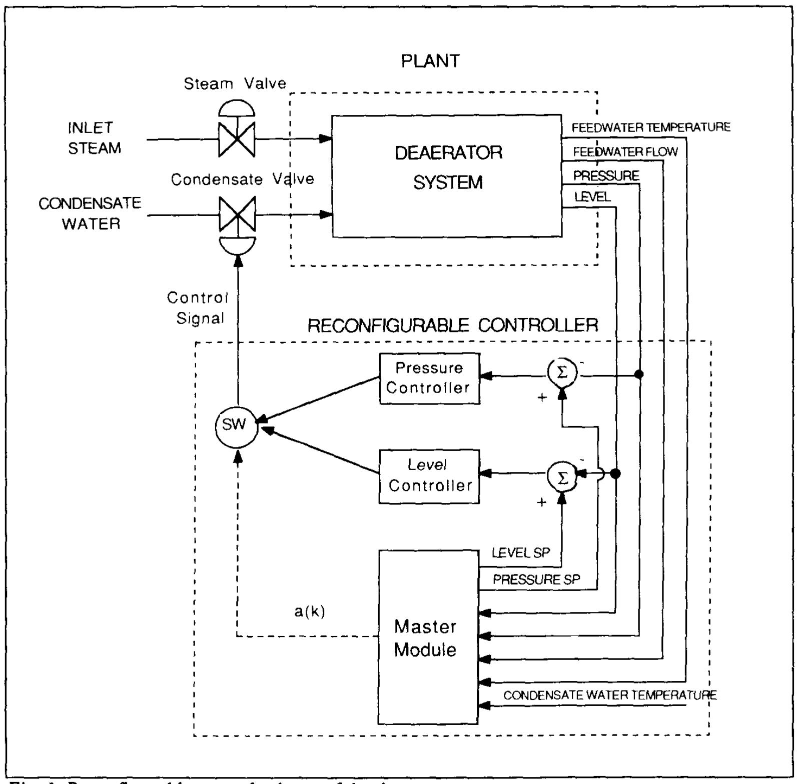

Fig. 1. Reconfigurable control scheme of the deaerator.

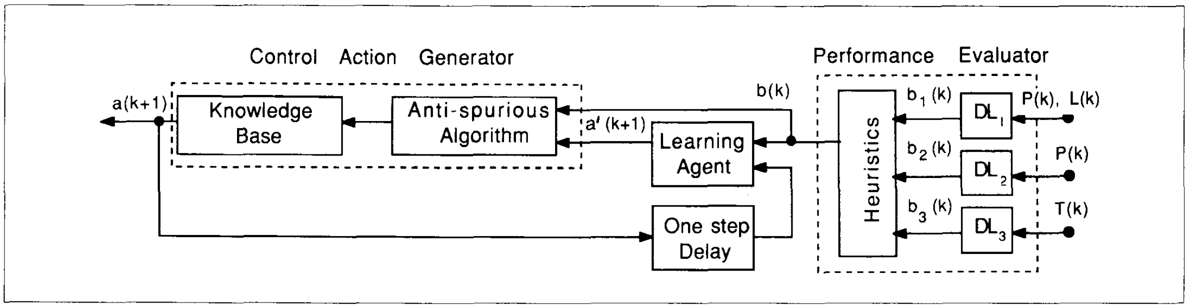

Fig. 2. Internal structure of the master module.

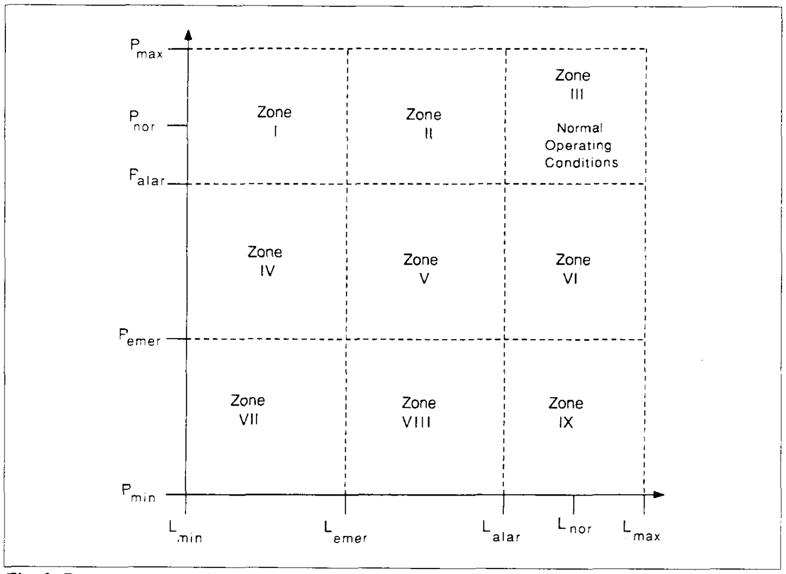

Fig. 3. Deaerator operating zones.

pressure is not a manipulated variable; instead a change in its position causes a disturbance in the simulated plant. In addition to generating the action selection signal $a(k)$, the master module is responsible for updating the set point of the selected controller if it is required under the current operating mode. Based on a predefined response criterion, the selected controller is expected to perform better than the other controllers for the given plant conditions. The master module incorporates learning mechanisms to gain sufficient knowledge for identification of the correct controller. In this design, the master module is composed of the three submodules: i) performance evaluator, ii) learning agent, and iii) control action generator, as shown in Fig. 2.

Performance Evaluator

On the basis of the measured plant response, the performance evaluator interprets the current performance of the selected controller as $b(k)$ having exactly one of the three values -1, 0, and 1. If the performance criterion is satisfied, a reward condition is indicated by setting $b(k)$ to 0; otherwise, a penalty condition $b(k) = -1$ is imposed. An inaction condition $b(k) = 1$ is asserted if the performance evaluator cannot arrive at a conclusion about the controller's performance. Three different rule-based procedures are used for generating $b(k)$ where the resulting performance indices are denoted by b_1, b_2, and b_3 (each having exactly one of the three values -1, 0, and 1). The measurement history (including the current values) of pressure and level is used for computing b_1, the pressure decay rate for b_2, and temperatures of both the condensate and the feedwater flows for b_3. A brief description of each of these b_i's follows.

Index b_1 identifies the operational zone of the plant at each sampling instant. Fig. 3 shows the partitioning of the entire operating regime of the deaerator into nine mutually exclusive but exhaustive zones such that the plant occupies exactly one of these zones at any instant of time. The objective is to select the appropriate controller on the basis of the information of the plant operating zone. For example, under the zones V, VI, VIII, or IX in Fig. 3, the pressure controller would be more suitable than the level controller because the level is above the emergency value but the pressure is under its alarm value. In contrast, for zones I, IV, or VII where the level is lower than its specified emergency value, level control is more important to avoid running the deaerator empty. Finally, in zone II and under the normal operating conditions in zone III, the level controller would be the appropriate choice.

The performance decision is also made dependent on the history of the level and the pressure measurements. The idea is to evaluate the performances of both controllers on the basis of the current operating condition and predicted plant dynamics. For example, in zone V, if the pressure is recovering (even though it is below the alarm setting P_{alar}) but the level is approaching L_{emer}, then the level controller should be selected. The pressure controller would be the appropriate choice in the reverse situation. Generally speaking, if a process variable (e.g., level or pressure) experiences recovery while the other one is degrading, there is a potential of switching controllers.

Index b_2 is intended to reduce the pressure decay rate so that the potential problem of cavitation at the feedwater pump can be avoided. To this effect, the deaerator pressure $P(k)$ is periodically measured and a weighted average $\delta P(k)$ of the pressure decay rate is obtained as follows:

$$\delta P(k) = \frac{\sum_{j=0}^{n(t)-1} \nabla P\,(k-j)\,\omega\,(j)}{T}$$

where the difference operator ∇ is defined as $\nabla f(k) = f(k) - f(k-1)$, the weight $w(j)$ is appropriately chosen (a possible choice is $w(j) = (j+1)^{-1}$), and the number n of past values of pressure readings over which the weighted average is taken is inversely related to the sampling interval T. Then, $\delta P(k)$ is compared with an *a priori* specified maximum pressure decay rate, δP_{max}, which is a function of the feedwater flow. If $\delta P(k) > \delta P_{max}$, the pressure controller is the appropriate action. An alarm signal should be activiated in this situation as a warning for a rapid pressure decay rate. Whenever $\delta P(k) < \delta P_{max}$, the pressure alarm is deactivated and b_2 set to -1 as it does not generate any decisive information.

Index b_3 bases the controller performance evaluation by taking into account the effects of the temperatures, T_{con} and T_{fw} of the condensate and feedwater, respectively. Since $T_{con} < T_{fw}$ under normal circumstances, a reduction of the condensate flow into the deaerator would be appropriate for mitigating the pressure loss, even though it may be contrary to maintaining the water level. However, pressure control is not the correct choice under unusual conditions for which $T_{con} > T_{fw}$. The rationale is that level control, under these circumstances, would increase both level and pressure by augmenting the condensate flow. This temperature-dependent criterion is implemented by setting b_3 to 1 if the pressure controller is active and $T_{con} > T_{fw}$; otherwise, b_3 is set to 0.

At each sampling instant k, the signal b is heuristically generated as a combination of b_1, b_2 and b_3 as follows:

$b = b_3$, if $b_3 = 1$, and any b_1 and b_2;
$b = b_2$, if $b_3 \neq 1$, $b_2 \neq 1$, and any b_1;
$b = b_1$, if $b_3 \neq 1$, and $b_2 = -1$.

The signal $b(k)$ is an input to both the learning agent and the control action generator. If the performance evaluator cannot make a decision on the current controller's performance, i.e., $b(k) = -1$, the learning mechanism is "frozen", i.e., the states of the automaton are held constant, and any switching from one controller to another is inhibited during the kth sampling period. On the other hand, if the performance evaluator does make a decision, i.e., $b(k) = 0$ or 1, then not only the states of the learning automaton are allowed to be updated but also $b(k)$ serves as the forcing function in the difference equations that model the learning automaton. For $b(k) = 0$ or 1, the performance evaluator has no direct influence on the control action generator.

Learning Agent

The learning agent is implemented using the concept of stochastic automata. At the sampling instant k, a specific controller is proposed for the $(k+1)$st sampling period. The inputs to the learning agent are the decisive response of the performance evaluator, i.e., $b(k)$ being 0 or 1, and $a(k)$ from the control action generator serves as the identity of the current controller. The output $a'(k+1)$ is the proposed choice for controller selection because the final decision, $a(k+1)$, is made by the control action generator. As discussed in the next section, $a(k)$ may differ from $a'(k)$.

Dynamic characteristics of the learning agent are essentially determined by its reinforcement scheme which is modeled as a set of r coupled difference equations where r = 2 in this example) is the number of alternative controllers. Several different types of reinforcement schemes have been reported in the literature [2]-[4],[10]. If a very high reliability is desired, several learning automata can be grouped together to form a learning module [3],[4] that may operate on majority voting.

Control Action Generator

The control action generator, as shown in Fig. 2, decides which one of the available controllers should act upon the plant. The control action generator governs switchings from one controller to another based on the plant response. As pointed out earlier, any such switching is inhibited during the kth sampling period if the performance evaluator fails to arrive at a decision, i.e., $b(k) = -1$.

The control action generator uses a filter to suppress the spurious fluctuations in $a'(k)$ such that the probability of erratic transitions between the controllers is reduced. These fluctuations may occur due to measurement noise, plant disturbances, and modeling uncertainties of the reinforcement scheme. A procedure, based on sequential testing such as [11], rejects spurious changes from one controller to another, and any such changes are accepted only if they are persistent.

The control action generator invokes certain constraints in the learning behavior of the master module. The decision $a'(k)$, proposed

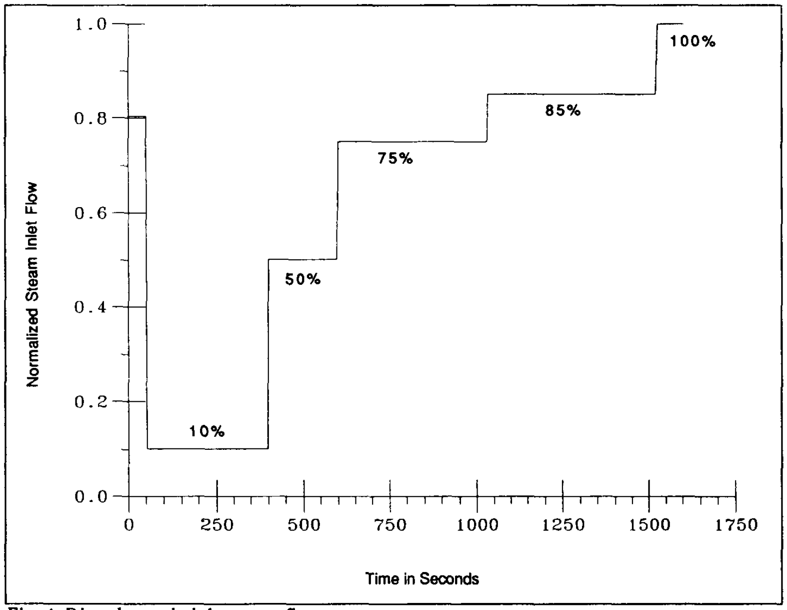

Fig. 4. Disturbance in inlet steam flow.

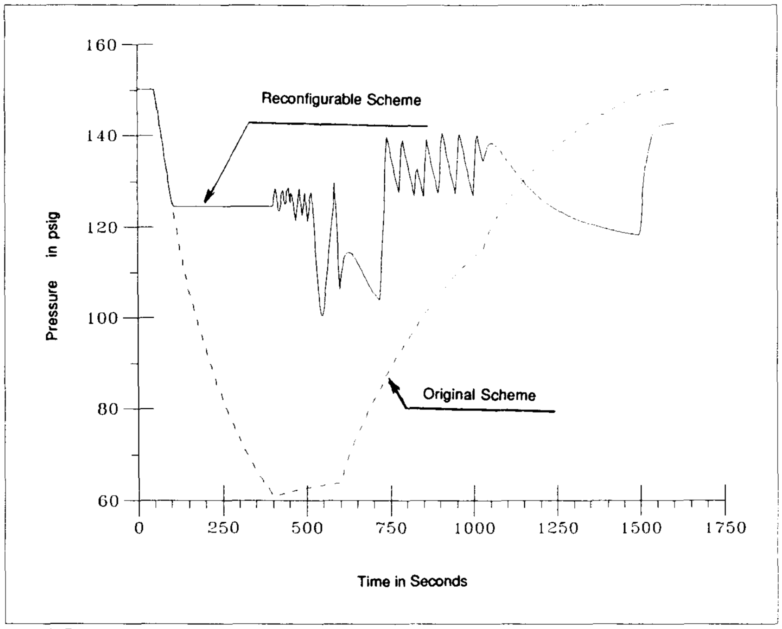

Fig. 5. Dearator pressure transients.

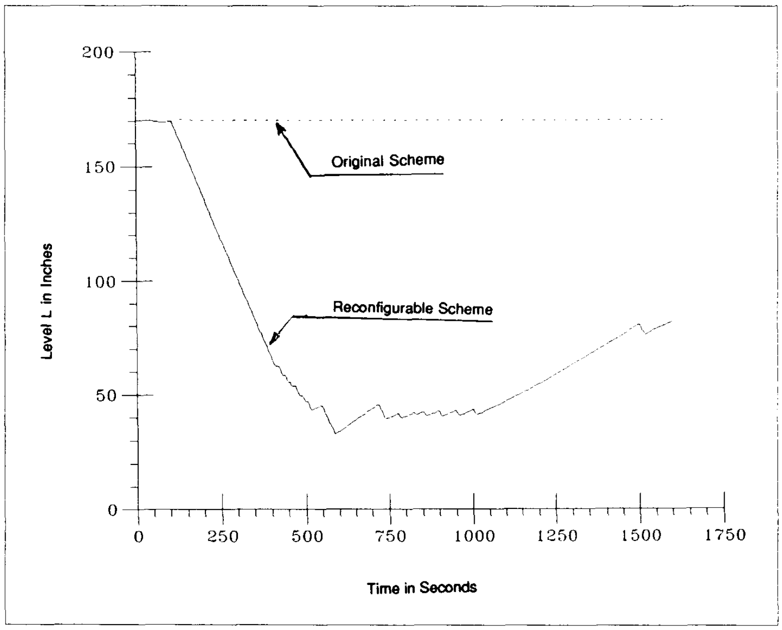

Fig. 6. Deaerator level transients.

specifications because it is unaware of the plant operational constraints. Whenever any such constraint is violated, the proposed action $a'(k)$ is overridden by the control action generator according to certain rules, and the next "best" choice is accepted as the action $a(k)$ for controller selection. These guidelines are to be derived from plant technical specifications, limitations of the individual controllers, operators' experience, and the qualitative knowledge about the plant and controller characteristics. If these guidelines can be represented as a set of rules, a real-time expert system would be an appropriate choice for making the final decision on controller selection. This is an area of future research and has not been addressed in this article.

Simulation Results and Discussion

The reconfigurable control scheme, described above, has been evaluated via interactive simulation using the Modular Modeling System (MMS) [12] and the Advanced Continuous Simulation Language (ACSL) [13]. The objective of the simulation experiments is to show that the learning approach provides a mechanism for adapting to a wide variety of operational scenarios including those that are not anticipated.

The sampling period, T, was set to 1 s and the maximum admissible pressure decay rate to 0.5 psi/s. These settings are compatible with the process dynamics of the deaerator. The Discrete Linear Reward Penalty scheme (DL_{RP}) [10] was chosen as the reinforcement scheme in the learning agent with the number of steps being equal to 6. The maximum, normal, alarm, emergency and minimum limits in Fig. 3 were chosen to be 180, 150, 125, 100, and 0 psig, respectively, for pressure, and 240, 170, 80, 50, and 10 in, respectively, for level. One of the various scenarios that were simulated to test the reconfigurable controller, is an abrupt reduction of the steam inlet flow into the deaerator followed by a gradual recuperation as shown in Fig. 4. The responses of the pressure $P(t)$ and level $L(t)$ are presented in Figs. 5 and 6, respectively, for the original (i.e., without learning) and the reconfigurable schemes.

As mentioned earlier, the response of the deaerator during a loss of steam flow is expected to be a reduction in pressure to the point that the plant must be shutdown to prevent damage to the feedwater pumps due to lack of available suction head. With the normal level controller still in operation, the flow of relatively cool condensate is increased to makeup for the loss in mass flow from the steam. However, this control action to maintain level at the set point further aggravates the pressure decay rate as seen in Fig. 5. Under these circumstances, switching to the alternate pressure controller would reduce the flow of cool condensate so that the deaerator pressure can be maintained while allowing the level to drop as seen in Fig. 6. The learning system, instead of having a preplanned schedule, decides the appropriate control action by monitoring of the process variables. The aspects of on-line monitoring, learning, and decision-making of the reconfigurable controller are discussed below.

Under normal steady operations, the steam valve is 80% open and the reconfigurable controller selects the level controller without switching to the pressure controller. Steam flow was stepped down to 10% at $t = 50$ s. Up to $t = 110$ s, i.e., 60 s after the disturbance in steam inlet flow, the reconfigurable control scheme continues to use the level controller. At this instant, the reconfigurable scheme identifies the average pressure decay rate in excess of 0.5 psi/s, and decides to switch to the pressure controller. As seen in Figs. 5 and 6, this switching arrests the pressure decay at the expense of letting the level drop. Under the level controller in the original scheme, the pressure continues to fall at an approximate rate of 0.5 psi/s. The reconfigurable scheme maintains $P(t)$ constant at 125 psig until a partial recuperation of the steam valve to 50% open occurs at $t = 400$ s. In contrast, under the original scheme, the pressure monotonically decreases, which would force the plant operator to take additional measures to protect the feedwater pumps. For example, if the pressure is allowed to drop below the emergency value, the plant would be forced to be shut down.

With the partial recuperation of the steam flow at $t = 400$ s, the reconfigurable scheme has to decide which one of the two controllers is more appropriate under the new environment. This results in switching back and forth between the two controllers to prevent the deaerator from running dry while keeping both pressure and its decay rate within the tolerable ranges. These switchings between controllers can cause oscillations of the processes variables $P(t)$ and $L(t)$ as seen in Figs. 5 and 6. The oscillations are more pronunciated in $P(t)$ than in $L(t)$ because of the faster dynamics of $P(t)$. With an additional recuperation of the steam supply to 75% at $t = 600$ s, the switchings become less frequent and so do the oscillations because pressure variations are less sensitive to changes in the incoming flow of the cool condensate water. At $t = 1050$ s, when the steam inlet valve is set to 85%, the rate of energy input into the deaerator is high enough so that the reconfigurable scheme can hold on to the level controller for a longer period. The resulting pressure and its decay rate are within tolerable limits, while the water level gradually recovers. Finally, at $t = 1500$ s, the steam inlet valve is brought to the 100% and the reconfigurable scheme switches to the level controller and remains there.

The phenomenon of more frequent switching at smaller steam inlet flow can be explained as follows. Whenever a switching from the pressure controller to the level controller takes place, the condensate flow is increased from a minimum to an almost maximum. These variations in operating conditions occur due to the reverse effects that the two controllers exert over the plant. More sophisticated algorithms and a larger number of alternative control actions would potentially reduce the frequency of switchings and thereby improve the performance of the reconfigurable scheme. For example, an advanced multi-input control law such as LQG [6],[7] would simultaneously minimize errors in the critical process variables so that the level is more effectively recuperated while maintaining the pressure within the safety ranges.

The performance of the reconfigurable scheme in a noisy environment has also been evaluated. We assume that the measured variables (i.e., level, pressure, condensate water temperature, and feedwater flow and temperature) used in the reconfigurable control system in Fig. 1, are subjected to a zero-mean additive noise with a uniform distribution between ±5% of the respective maximum ranges listed in Table II. Figs. 7 and 8 present the responses of the pressure $P(t)$ and level $L(t)$, respectively, for the original and the reconfigurable schemes. Both the noise-contaminated and expected values of $L(t)$ and $P(t)$ are plotted to show that they are maintained well above their respective emergency limits, even though the available information is corrupted. These results show that performance of the reconfigurable scheme suffers from noise contamination but it is still satisfactory and significantly superior to that of the original

Table II: Ranges of Measured Variables

Sensor Description	Maximum Reading	Minimum Reading	Range
Water Level, L (in)	300	0	300
Pressure, P (psig)	300	0	300
Condensate Temp, T_{con} (°F)	400	0	400
Feedwater Temp, T_{fw}(°F)	400	0	400
Feedwater Flow, W_{fw} (lbm/h)	$0.30(10^6)$	$0.05(10^6)$	$0.25(10^6)$

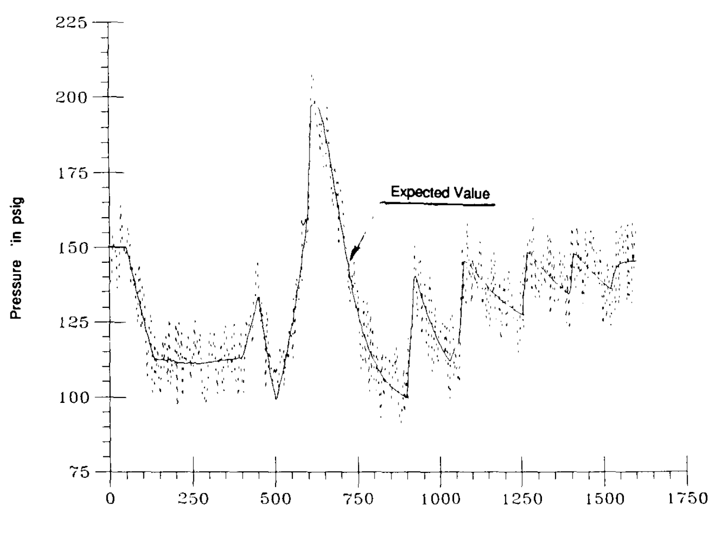

Fig. 7. Dearator pressure transients using the reconfigurable scheme (with noisy measurements).

scheme without any learning capability. Thus, the stochastic behavior of the learning automata proves to be effective in uncertain environments.

Summary and Conclusions

A reconfigurable control scheme is proposed which, unlike a preprogrammed one, uses stochastic automata to learn the current operating status of the plant by dynamically monitoring the system performance and then switching to the appropriate controller on the basis of the observed performance. The potential applicability of this reconfigurable control scheme to electric power plants has been investigated. A deaerating feedwater heater, equipped with a water level controller and a pressure controller, has been chosen to study the feasibility of reconfigurable control in power plant operations. Simulation experiments have been conducted on the basis of a model of the Experimental Breeder Reactor (EBR-II) at the Argonne National Laboratory site in Idaho. The results show that the reconfigurable control scheme is capable of providing a sufficient margin for the net positive suction head at the feedwater pumps under loss of steam flow into the deaerator. Under similar circumstances, the existing controller in the deaerator would be incapable of maintaining the pressure and its decay rate within the safe margins, and thereby force the plant operator to take additional measures to protect the feedwater pumps. The learning agent in the reconfigurable controller of the deaerator is capable of taking a correct action in the (unusual) event of the condensate water being warmer than the feed water. This example shows how the control system can learn to react to unanticipated circumstamces which could be difficult to be handled by human operators within time constraints.

Incorporation of learning capabilities within the reconfigurable control scheme is promising for unanticipated and uncertain plant conditions for which preprogrammed control algorithms are apparently difficult to formulate. Both the performance evaluator and the set of alternative controllers are critical for the reconfigurable control scheme, and may rely on a combination of analytical and heuristic techniques. While the general structure of the performance evaluator is specifically dependent on the given application, the individual control algorithms are more likely to be formulated by taking advantage of the existing model-based [6,7] and rule-based [8] methodologies.

The reconfigurable level/pressure control system, presented in this article, is an illustration of how learning automata can be applied to continuous processes like power plants. Further research is required for investigating the feasibility of the proposed concept for control of continuous processes under diverse operating conditions. One potential area of research is automated start-up of fossil and nuclear power plants.

As a power plant traverses through a wide range of operating conditions from start-up to full power, the control functions are reconfigured according to a prescribed procedure. However, deviations from the normal course of actions, which are very likely to occur due to uncertainties such as those imposed by overhaul of plant components and recalibration of instruments, would slow down the start-up process and run the risk of damaging plant component(s). The goal is to reduce the start-up time while avoiding any damage to the plant. This could be accomplished by reconfigurable control, based on learning automata, which would identify appropriate actions in real time whenever there is a mismatch between the observed symptoms and prescribed conditions. The learning automata do not replace the plant protection and safety systems but serve as a potential means to steer the plant away from undesirable operating conditions. This capability is derived, on the basis of the observed plant performance, by selecting the most appropriate control action.

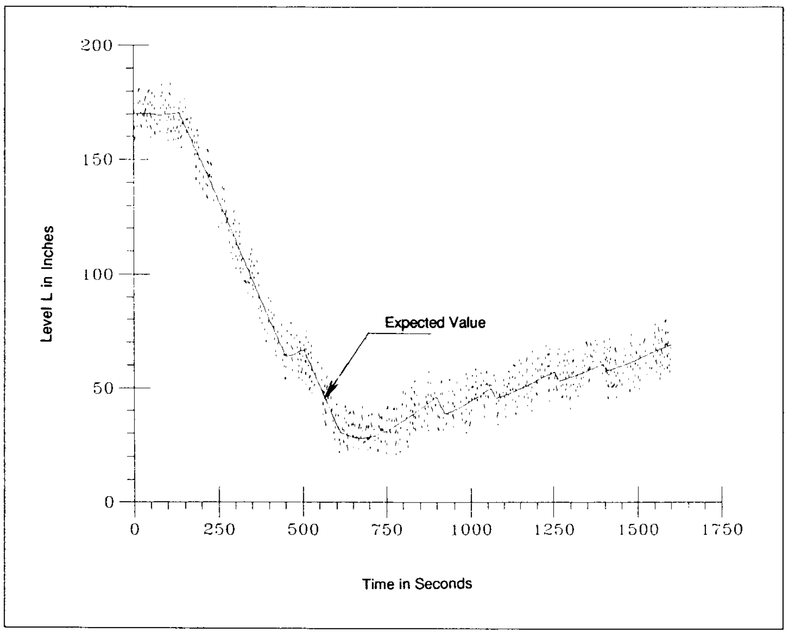

Fig. 8. Deaerator level transients using the reconfigurable scheme (with noisy measurements).

Acknowledgment

The authors gratefully acknowledge comments and suggestions of Professors E.H. Klevans and D.R. Mudgett.

References

[1] M. Itoh, I.Tai, K. Monta and K. Sekimizu, "Artificial Intelligence Applications for Operation and Maintenance: Japanese Industry Cases", *Artificial Intelligence and Other Innovative Applications in the Nuclear Industry*, (eds.) M. C. Majumdar, D. Majumdar, and J. I. Sackett, Plenum Press, pp. 37-44, 1988.

[2] K. Narendra and M.A.L. Thathachar, "Learning Automata- A Survey," *IEEE Trans. Sys., Man, Cybern.*, Vol. SMC-4, No. 4, pp. 323-334, July 1974.

[3] K. Narendra and M.A.L. Thathachar, *Learning Automata: An Introduction*, Prentice-Hall, 1989.

[4] H.E. Garcia, "Learning Automata in a Reconfigurable Controller," M.S.E.E. Thesis, The Pennsylvania State University, December 1989.

[5] R.S. Ornedo *et al.*, "Design and experimental evaluation of an automatically reconfigurable controller for process plants," in *Proc. American Control Conference*, Minneapolis, pp. 1662-1668, June 1987.

[6] B.D.O. Amderson and J.B. Moore, *Optimal Control: Linear Quadratic Methods*, Prentice Hall, Englewood Cliffs, NJ, 1990.

[7] J.P. McDonald and H.G. Kwatny, "Design and Analysis of Boiler-Turbine-Generator Controls Using Optimal Linear Regulator Theory," *IEEE Trans. Aut. Contr.*, vol. AC-18, no. 3, pp. 202-209, June 1973.

[8] J.A. Bernard, "Use of a Rule-Based System for Process Control," *IEEE Control Systems Magazine*, vol. 8, no. 5, pp. 3-13, October 1989.

[9] A. Ray, "Transients in a Power-plant Feed-pump System," *Simulation Council Proceedings Series. Simulation of Energy Systems Part 2*, ed. K.E.F. Watt, Chapter 12, December 1978.

[10] B.J. Oommen and J.P. Christensen, "ε-Optimal Discretized Linear Reward-Penalty Learning Automata," *IEEE Trans. Sys., Man, Cybern.*, vol. SMC-18, no. 3, pp. 451-458, May/June 1988.

[11] A. Ray, "Fault Detection in Multiply-Redundant Measurement Systems via Sequential Testing," *ASME Journal of Dynamic Systems, Measurement, and Control*, pp. 329-332, June 1989.

[12] Modular Modeling System (MMS): A Code for Dynamic Simulation of Fossil and Nuclear Power Plants, Overview and General Theory, *EPRI Report No. CS/NP-2989*, Electric Power Research Institute, Palo Alto, CA, March 1983.

[13] ACSL Reference Manual, 1987, Mitchell and Gauthier Associates, 73 Junction Square Drive, Concord, Massachusetts 01742.

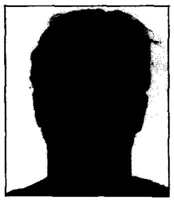

Humberto E. Garcia received the Electrical Engineering Diploma from the Universidad de Carabobo in 1985 and the M.S. degree in Electrical Engineering from the Pennsylvania State University in 1989. He is currently pursuing the Ph.D. degree in Electrical Engineering at Penn State. His current research interests include distributed intelligent systems, intelligent control systems, computerized automation and control, decision support systems, and artificial intelligence. Mr. Garcia is a member of Tao Beta Pi.

Asok Ray earned a Ph.D. degree in Mechanical Engineering and graduate degrees in each of Electrical Engineering, Computer Science, and Mathematics. Dr. Ray has held research and management positions at GTE Strategic Systems Division, Charles Stark Draper Laboratory, and MITRE Corporation as well as academic positions at Carnegie-Mellon University and Massachussetts Institute of Technology. Dr. Ray joined the Pennsylvania State University in July 1985, and is currently a Professor of Mechanical Engineering. Dr. Ray's research experience includes real-time microcomputer-based control and instrumentation, networking and communication protocols, intelligent systems design, and modeling and simulation of dynamical systems as applied to aeronautics, process control, and autonomous manufacturing. Current research interests of Dr. Ray include distributed control systems, applied stochastic processes and fault detection, and intelligent instrumentation and computer networking for aeronautical and manufacturing systems. Dr. Ray has authored or co-authored over one hundred and fifty research publications, and is an Associate Fellow of AIAA, a Senior Member of IEEE and a member of ASME.

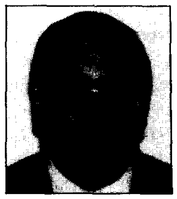

Robert M. Edwards received the B.S. degree in Nuclear Engineering from the Pennsylvania State University in 1971 and the M.S. degree in Nuclear Engineering from the University of Wisconsin in 1972. His professional experience includes power plant controls, modeling and transient analysis at General Atomic in San Diego, California during the early 70s. Before returning to Penn State to pursue the Ph.D. degree in Nuclear Engineering in 1986, he was the director of software development for a manufacturer of microprocessor based image analysis systems. At Penn State he is employed as a full time research assistant while nearing the completion of his Ph.D. thesis and has conducted research on advanced controls and diagnostics sponsored by the Department of Energy, National Science Foundation, and Argonne National Laboratory. Mr. Edwards is a Registered Professional Engineer (Mechanical) in California and is also a member of the American Nuclear Society and Society for Computer Simulation.