Theory and Application of an Extended Horizon Self-Tuning Controller

A robust version of the self-tuning regulator is developed. The regulator, which requires relatively little knowledge of system characteristics (estimated order of transfer function polynomials and an upper bound for transportation delays), has been shown to yield stable control and convergence for linear, time-invariant systems. Simulations and practical tests on a large pilot-scale process have shown that the inclusion of a variable forgetting factor and an "extended horizon" control criterion provides the regulator with a sufficient degree of robustness and flexibility to perform well in a nonlinear time-varying environment. The regulator makes use of intuitively easy-to-understand concepts and leaves few degrees of freedom for the potential user. Furthermore, extensive experiments and simulation studies have shown it to be insensitive to choice of initial conditions and dynamic characteristics set by the user.

B. E. YDSTIE,
L. S. KERSHENBAUM,
and R. W. H. SARGENT
Department of Chemical Engineering
and Chemical Technology
Imperial College of Science and Technology
London SW7 2BY England

SCOPE

The idea of the self-tuning regulator was developed (Åström and Wittenmark, 1973) to cope with the problem of controlling linear, time-invariant systems with unknown or uncertain models, and which were subject to stochastic disturbances—a situation which is not common in the chemical industry. There have been subsequent attempts to modify the basic concepts to obtain algorithms which can deal with the more typical chemical engineering problem—the control of processes with nonlinear, uncertain, nonconstant dynamics and with long and variable transportation delays.

In order to deal with this more complex problem it is necessary to incorporate some adaptive mechanism into the part of the algorithm which estimates a model for the process being controlled. It is also necessary to "detune" the control action in an appropriate way to prevent severe overshoot and potential instability of the real nonlinear process. Finally, it is important to reduce the sensitivity of the algorithm to any remaining parameters which must be specified a priori by the operator or control engineer.

Only if all the above are achieved is the resulting control algorithm likely to be robust enough to be generally useful in the control of process plants.

CONCLUSIONS AND SIGNIFICANCE

It is likely that some form of self-tuning controller will find significant application in the control of time-dependent nonlinear processes within the chemical industry. These controllers have the advantages of simplicity, flexibility, and compatability with microcomputer technology, as well as an intuitively satisfying basis.

The results of this study have shown that the basic self-tuning algorithm can be modified to produce a version far more robust than the early formulations. The algorithm has been shown to be stable and convergent for linear, time-invariant systems, and requires only a few parameters to be supplied by the operator. Moreover, for a wide range of simulations and experiments on a large-scale process plant, the algorithm performs extremely well in nonlinear, time-dependent situations and is relatively insensitive to the actual numerical values of the few parameters which do have to be chosen by the user. Specifically, knowledge of the time delay in the system and the order of the linearized system model (both of which are difficult to predict a priori and are generally dependent upon the local operating conditions) are not required for the practical implementation of the algorithm.

It is felt that further developments in this area could lead to a general-purpose control algorithm of sufficient flexibility and robustness to gain wide usage in the chemical industry.

Correspondence concerning this paper should be addressed to L. S. Kershenbaum.
B. E. Ydstie is presently at Chemical Engineering Department, University of Massachusetts, Amherst, MA 01003.

INTRODUCTION

Accurate dynamic modeling of chemical engineering systems is by no means a trivial task. The processes in question are generally nonlinear, time-varying, and have variable time delays and other nonminimum phase effects. In addition, they are often interactive and are perturbed by changes in feedstock and operating conditions. All of these factors make dynamic modeling and control somewhat different from that of typical electrical and mechanical systems where linear, time-invariant stochastic models often give sufficient flexibility and accuracy even in a high performance environment. Accordingly, application of advanced control strategies in the chemical industries has had only limited success because it often involves a laborious and costly exercise to develop accurate process models and advanced control devices. The single-input single-output Ziegler-Nichols tuned three-term regulator is still the most successful industrial regulator, and advances made in control theory during the past two decades have thus made little impact in the industrial community.

However, the picture outlined above is now changing. Cheap and powerful microcomputers are available at low cost and in situ data filtering and parameter optimization need no longer be procedures reserved for those who are willing to invest in large computers. Moreover, recent advances made in the theory of robust adaptive prediction and control following from the pioneering works of Peterka (1970) and Åström and Wittenmark (1973) point toward the possibility of wide application of advanced control in industry. Self-tuning regulators have by now been used successfully to control and optimize a wide variety of process plants such as paper machines and ore crushers (Åström et al., 1977), batch chemical reactors and effluent pH treatment systems (Clarke and Gawthrop, 1981), cement kilns (Dumont and Bélanger, 1981), and chip refiners (Dumont, 1982). All of these applications use relatively simple theory for estimator and controller design. Seborg et al. (1983) give many more applications and a review of current theory.

The scope for developing new regulators and more robust strategies is almost endless, and many of these have been reported in the literature. However, the gap between theory and application still exists and often these strategies offer little intuitive guidance for a practicing engineer on how to implement and initialize the adaptive strategy. Although the schemes suggest that the regulators adjust themselves, there is invariably a set of parameters that must be set by the user, and good performance is often critically dependent on "correct" choice of these.

There are also special points that need to be observed when we apply adaptive theory to industrial processes. These are concerned with integrity, robustness, and stability characteristics when the regulators are subjected to adverse conditions such as long periods of operation during which little or no dynamic information about the plant is available for lack of excitation. It is also important to bear in mind that most of the theory is, in essence, linear and time-invariant, and therefore not directly applicable to industrial situations. Nevertheless, insight can be brought forth through linear analysis and it appears important to justify adaptive schemes not only through experiments, but also through rigorous stability and convergence theory.

In this paper we explore one set of adaptive schemes in detail, referring to theoretical and simulation results as well as to a wide range of tests performed on a large pilot-scale process plant. To enable the regulator to cope more easily with the special problems met in industrial control, we have modified the original self-tuning scheme. The most significant modifications are: past data are discounted by the variable forgetting factor approach of Fortescue et al. (1981); the minimum variance control objective is replaced by the extended horizon controller of Ydstie (1982) which detunes performance and obviates the need for knowing the system time delay; the algorithm is expressed in the incremental form of Fortescue (1977) which eliminates offset.

Problem Formulation

We assume in the following that, whatever the actual form of the process model, the local behavior of the process can be adequately described by a linear, finite-dimensional, input-output model. The parameters of this model will be estimated on-line from available measurements. We are seeking an adaptive controller which yields stable input-output behavior during transients and that asymptotically achieves optimal performance for a linear, time-invariant system.

It is evident that if an approximate linear model is to be used successfully in a time-varying nonlinear environment, it is necessary to discount old data to enable the estimator to adapt to changes and upsets in the process. One popular way of giving more weight to recent data is to incorporate a constant forgetting factor less than unity in a recursive estimator. This strategy may work well for noisy plants, but on near-deterministic systems, which are typical of chemical plants, a constant forgetting factor will lead to "covariance wind-up" and estimator instability. Fortescue et al. (1981) discussed this problem in some detail and showed that this instability is due to loss of excitation of the system. One source of the problem is that useful information is continually discarded because the variance of past data is artificially increased even when there is little or no new dynamic information coming in from the plant. Eventually the plant becomes temporarily or permanently unstable depending on the state of the estimator and nonlinearities in the process. In that work it was suggested to define a measure of the information content of the estimator and to choose a forgetting factor at each step to maintain this measure constant. Covariance wind-up is thereby avoided while adaptivity is retained in such a way that the process can adapt itself to both slow and sudden changes in process dynamics.

The scheme, as outlined by Fortescue, has already been tried successfully on a chip refiner (Dumont, 1982); it is in continuous operation on a commercial paper machine control scheme (Fjeld et al., 1981), a Kalman filter application (Doss et al., 1983), and a river water level control system (Fjeld et al., 1983). Further studies on the pilot plants at Imperial College suggest that the variable forgetting factor approach is indeed a robust and versatile method of discounting past information. The scheme has recently been developed further, and we give here a more exact version of the update formula that requires fewer parameters to be set and that provides better intuitive guidance on how the information measure should be chosen.

The early works on self-tuning regulators used a minimum variance regulator which, in turn, required knowledge of the system delay time. However, several applications on our own plants (Ahmad, 1981; Ydstie, 1982) demonstrated that this control objective is often too strong and may lead to drastic control action which in turn produces overshoot, intersample oscillations, or complete instability when the process is not strictly minimum-phase or when the delay time is incorrectly specified. Nonminimum phase (inverse response) characteristics, however, appear to be common in modeling chemical processes; this weakness together with the fact that process dead times are neither generally known a priori nor constant, motivated a search for alternative control strategies. Clarke and Gawthrop (1975, 1979) introduced a modified control objective in which both control action and output error are weighted in the cost function; the later version included three user-defined transfer functions, which effectively allows the user to prescribe the closed-loop response. However the relation between the adjustable parameters and the form of the response is indirect, and more recent developments (Wellstead et al., 1979, Åström and Wittenmark, 1980; Goodwin and Sin, 1981) use pole-zero assignment rather than optimal criteria in the controller design. The algorithms are complex, and considerable experience and insight is required in making an appropriate choice of parameters.

A simpler approach is to choose the control to follow a prescribed recovery trajectory (typically defined as a first-order response). This is related to the Clarke-Gawthrop (1979) controller, the Model Algorithmic Control (MAC), first proposed by Richalet et al. (1978) and reviewed by Mehra et al. (1982), and is also the basis of the Dynamic Matrix Control (DMC) algorithm proposed by Cutler and Ramaker (1980); both of these latter algorithms are based on linear convolution models rather than the ARMA model traditionally used for self-tuning controllers. Adaptive versions of these types of controllers have been suggested by Tung (1983) and Lee and Lee (1983).

The accurate profiling of the response is not really an important issue, and the essence of all these approaches is to desensitize, or detune, the controller in order to achieve robustness. In this work we propose a detuning technique which maintains both the intuitive and computational simplicity of the original self-tuning idea. Instead of driving the output error to zero in one step (that is, within the natural process delay time), we simply allow the controller more time. By extending the controller's time horizon in this way, it can "look beyond" time delays and periods of inverse response, thus rendering it unnecessary to predict the natural time delay, and making it possible to deal with variable time delays and other nonminimum phase effects. The user has only to choose the time horizon to control the amount of detuning, and since this is related to the process response time he will have an intuitive idea of the required value even before the controller is put on line. In the original self-tuning regulator, the user was required to estimate the natural delay time, so the number of user-defined parameters is not increased, and indeed the choice of time horizon is much less critical in the case of the extended horizon controller.

Extended horizon prediction and control has been suggested and used earlier (Peterka, 1970; Tung, 1983; Lee and Lee, 1983; Mosca and Zappa, 1984); however, the formulation which we propose is simpler computationally, has a more intuitive basis, and is directly related to the original self-tuning controller.

As already noted, the parameters in the process model are estimated on-line, and it is well known that for a nonlinear system the resulting output predictions may well be biased, leading to offset. It may therefore be advantageous to incorporate some kind of integral action in the controller to deal with this offset. The method suggested by Wittenmark (1973) and used by many other workers (Clarke and Gawthrop, 1975; Sastry et al., 1977) uses the difference between the output and the setpoint in place of the actual output in the control law. Unfortunately, this introduces an extra zero on the imaginary axis (on the unit circle for discrete formulations), which can seriously affect the stability and convergence of the controller. Fortescue (1977) suggests using an incremental model and regulator where this zero is exactly cancelled by a matching pole; thus no nonminimum phase effects are introduced by the integration. Simulation and plant experiments (Ydstie, 1982) show that this is indeed a more satisfactory way of reducing steady-state errors. Other methods to reduce offset are reviewed by Clarke et al. (1983).

THEORETICAL DEVELOPMENT

The local approximation to the (unknown) process model is assumed to be a linear discrete-time (ARMA) model of the form:

$$A(z^{-1})y_k = z^{-D}B(z^{-1})u_k + d + v_k \tag{1}$$

where z^{-1} is the unit delay operator ($z^{-1}\{y_k\} = y_{k-1}$), D is the estimate of the process delay time (≥ 1), d is a bias parameter, and v_k an independent noise sequence with zero mean and variance q_k. The polynomials A and B are written:

$$A(z^{-1}) \equiv 1 + a_1z^{-1} + a_2z^{-2} + \ldots + a_nz^{-n}$$
$$B(z^{-1}) \equiv b_o + b_1z^{-1} + b_2z^{-2} + \ldots + b_{n-1}z^{-n+1} \tag{2}$$

There is no loss of generality in assuming that polynomials A and B have the same number of unknown coefficients (which are to be estimated on-line), since trailing coefficients can be zero in either or both polynomials.

In the following we will first discuss possible control strategies for this model and then present suitable techniques for estimating the unknown parameters.

Extended Horizon Control

The minimum variance control law (Åström and Wittenmark, 1973) is obtained by choosing the control u_k so that $E\{y_{k+D}|Y_k\} = y^*_{k+D}$, where y^*_{k+D} is the desired output at time $(k + D)$ and Y_k denotes all the information available at time k. For the above control law, it is clear from Eqs. 1 and 2 that we must have the coefficient b_o well-defined and nonzero, and hence D must not be underestimated. Overestimation of D usually yields $b_o \neq 0$ and the control is usually satisfactory, but of course if the system is in fact linear, the assumed model cannot fit it and the control no longer yields minimum variance of the output error.

We prefer to provide flexibility by choosing an extended horizon control strategy in place of the minimum variance control law. The extended horizon control is chosen so that $E\{y_{k+T}|Y_k\} = y^*_{k+T}$ where the horizon T is greater than the true process delay time. One is then free to underestimate D (and hence obtain an exact fit of parameters if the process were linear) and, as will be shown below, still avoid the problem of near-zero critical coefficients of B. Furthermore, we can take explicit account of the fact that controls at times later than the current time k also affect the output y_{k+T}.

For the controller development it is convenient to express Eq. 1 in an equivalent prediction model form. Coefficients F and G can always be found which satisfy the polynomial identity:

$$1 \equiv F(z^{-1})A(z^{-1}) + z^{-T}G(z^{-1}) \tag{3}$$

where $F(z^{-1}) \equiv f_o + f_1z^{-1} + \ldots + f_{T-1}z^{-T+1}$ and $G(z^{-1}) \equiv g_o + g_1z^{-1} + \ldots + g_{n-1}z^{-n+1}$. Simple manipulations with Eqs. 1 and 3 yield:

$$(1 - z^{-T}G(z^{-1}))y_k = z^{-D}F(z^{-1})B(z^{-1})u_k + F(z^{-1})(v_k + d). \tag{4}$$

If the delay time were large and one were certain of its lower bound, one could choose D accordingly and be left with fewer parameters to be estimated on-line; in all of our work, however, we simply take $D = 1$, its minimum value, and the model can be expressed as:

$$y_k = \phi'_{k-T}\theta + \epsilon_k$$

or

$$y_{k+T} = \phi'_k\theta + \epsilon_{k+T} \tag{5}$$

where

$$\phi'_k \equiv (y_k, y_{k-1}, \ldots y_{k-n+1}; u_{k+T-1}, u_{k+T-2}, \ldots u_k, \ldots u_{k-n+1}; 1)$$

$$\theta' \equiv (\alpha_o, \alpha_1, \ldots \alpha_{n-1}; \beta_1, \beta_2, \ldots \beta_T, \ldots \beta_{n+T-1}; \delta)$$

$$\delta \equiv \sum_{i=o}^{T-1} f_i d; \quad \epsilon_k = \sum_{i=o}^{T-1} f_i v_{k-i}.$$

It thus follows that

$$E(\epsilon_k|Y_{k-T}) = 0; E(\epsilon_k^2|Y_{k-T}) = \sum_{i=o}^{T-1} f_i^2 q_{k-i} = r_k < \infty. \quad (6)$$

We have thus expressed the original model as a T-step ahead prediction model. We note in particular that if T is greater than the (unknown) true process delay time, then the output y_{k+T} depends upon future inputs ($u_{k+1}, u_{k+2}, \ldots$) as well as upon the current input, u_k. The extended horizon control criterion will thus be satisfied if controls $u_k, u_{k+1}, \ldots u_{k+T-1}$ are chosen so that

$$y_{k+T}^* - \phi_k'\theta = y_{k+T}^* - h_{k+T} - \sum_{i=1}^{T} \beta_i u_{k+T-i} = 0 \quad (7)$$

where

$$h_{k+T} \equiv \sum_{i=o}^{n-1} \alpha_i y_{k-i} + \sum_{i=1}^{n-1} \beta_{T+i} u_{k-i} + \delta \quad (8)$$

We see from Eq. 5 that h_{k+T} may be interpreted as the predicted value of y_{k+T} with $u_k = u_{k+1} \ldots = u_{k+T-1} = 0$, and hence it is known at time k.

There are many possible control strategies which satisfy Eq. 7, and we consider three obvious possibilities:

Strategy 1. Choose the constant control $u_k = u_{k+1} \ldots = u_{k+T-1}$ that satisfies Eq. 7. The law is then given by:

$$u_k = (y_{k+T}^* - h_{k+T}) \Big/ \sum_{j=1}^{T} \beta_j \quad (9)$$

Strategy 2. Choose the control u_k which together with $u_{k+1} = u_{k+2} = \ldots u_{k+T-1} = 0$ satisfies Eq. 7. This gives the law:

$$u_k = (y_{k+T}^* - h_{k+T})/\beta_T \quad (10)$$

Strategy 3. Choose the sequence of controls $u_k, u_{k+1}, \ldots u_{k+T-1}$ which satisfies Eq. 7 and minimizes the control effort, $\Sigma_{i=1}^{T} u_{k+T-i}^2$. The solution of this optimal control problem shows that the controls should be given by:

$$u_{k+T-i} = \beta_i (y_{k+T}^* - h_{k+T}) \Big/ \sum_{j=1}^{T} \beta_j^2; i = 1, \ldots, T$$

In the practical implementation we recompute the controls at each time-step and hence apply only the first control in the sequence, giving the law:

$$u_k = \beta_T (y_{k+T}^* - h_{k+T}) \Big/ \sum_{j=1}^{T} \beta_j^2 \quad (11)$$

(Indeed, even in the first two strategies only u_k is implemented and subsequent controls are recomputed at later time-steps).

If the horizon T is in fact equal to the true system delay time, any extended horizon strategy satisfying Eq. 7 will yield the minimum variance control law since $\beta_i = 0, i = 1, \ldots, T-1$. Larger values of T tend to reduce the controller gain, yielding the desired detuning.

In general, all these strategies give satisfactory controls. Since by definition at least one of the $\beta_j, j = 1, \ldots, T$, must be nonzero, the control called for by strategy 3 is always finite. There are certain linear systems which, for a specific value of T, have $\beta_T = 0$ and/or $\Sigma_{j=1}^{T}\beta_j = 0$; in these special cases, strategy 3 would apply no control and strategies 1 and 2 could call for infinite values of u_k. For a nonlinear system, the probability of such a situation occurring is negligible. Nevertheless, both strategies 1 and 2 can require very large values of u_k due to small values of $\Sigma_{j=1}^{T}\beta_j$ or β_T, respectively.

Strategy 3 appears to be the preferred choice among those considered and the one we have ourselves adopted recently; most of the results reported here, however, were obtained using strategy 1 and, in practice, the results are usually indistinguishable. There are numerous other possible strategies which may well be worth exploring.

Some closed-loop stability considerations for these controllers have been given by Ydstie (1984), Goodwin and Dugard (1983) and Ydstie and Liu (1984).

Integral Action

As noted earlier, it may be advantageous to include some form of integral action into the controller to deal with offset resulting from biased estimates.

The method developed by Fortescue (1977) has been used in all of our studies. It can be shown that new polynomials $\overline{F}$ and $\overline{G}$ can always be chosen to satisfy the modified polynomial identity

$$1 - z^{-T} = (1 - z^{-1})\overline{F}(z^{-1})A(z^{-1}) + (1 - z^{-1})z^{-T}\overline{G}(z^{-1}) \quad (12)$$

Manipulation of Eqs. 1 and 12 gives

$$y_k = y_{k-T} + \overline{G}(z^{-1})(1 - z^{-1})y_{k-T} + z^{-D}\overline{F}(z^{-1})B(z^{-1})(1 - z^{-1})u_k + \overline{F}(z^{-1})(1 - z^{-1})(d + v_k).$$

Noting that $(1 - z^{-1})y_k \equiv y_k - y_{k-1} \equiv \Delta y_k$, taking $D = 1$, and assuming a constant value of the bias parameter, d, we get the incremental model

$$y_{k+T} = y_k + \overline{\phi}_k'\overline{\theta} + \overline{\epsilon}_{k+T} \quad (13)$$

where

$$\overline{\phi}_k' \equiv (\Delta y_k, \Delta y_{k-1}, \ldots \Delta y_{k-n+1}; \Delta u_{k+T-1}, \ldots \Delta u_k, \ldots \Delta u_{k-n+1})$$

$$\overline{\theta}' \equiv (\overline{\alpha}_o, \overline{\alpha}_1, \ldots \overline{\alpha}_{n-1}; \overline{\beta}_1, \ldots \overline{\beta}_T, \ldots \overline{\beta}_{n+T-1}); \overline{\epsilon}_k \equiv \overline{F}(z^{-1})\Delta v_k$$

The corresponding control equations for our three strategies become:

$$\Delta u_k = (y_{k+T}^* - \overline{h}_{k+T}) \Big/ \sum_{j=1}^{T} \overline{\beta}_j \quad \text{for strategy 1} \quad (14)$$

$$\Delta u_k = (y_{k+T}^* - \overline{h}_{k+T})/\overline{\beta}_T \quad \text{for strategy 2} \quad (15)$$

$$\Delta u_k = \overline{\beta}_T(y_{k+T}^* - \overline{h}_{k+T}) \Big/ \sum_{j=1}^{T} \overline{\beta}_j^2 \quad \text{for strategy 3} \quad (16)$$

where

$$\overline{h}_{k+T} \equiv y_k + \sum_{i=o}^{n-1} \overline{\alpha}_i \Delta y_{k-i} + \sum_{i=1}^{n-1} \overline{\beta}_{T+i} \Delta u_{k-i}$$

(In later work, we have relaxed the assumption of constant d and allowed for the estimation of an additional parameter related to Δd; the performance remains unchanged.)

There are several advantages to the incremental scheme. Fortescue (1977) has shown that this formulation is equivalent to the original positional algorithm cascaded with an integrator to eliminate offset. Furthermore, he shows that the zero which is introduced onto the unit circle by the integration is exactly canceled by a matching pole; hence, no deterioration is produced in the stability, convergence, or dynamic performance of the algorithm. A second advantage of the incremental formulation is that it is directly compatible with those process control computers which operate in incremental mode.

Minimum Variance Estimation with Weighting of Data

All of the above assumes the existence of a model with known parameters. The parameters of the prediction model Eq. 5, are unknown and we require an estimation scheme that is robust with respect to changing process dynamics. Because the process we are approximating is nonlinear, any choice of method will be somewhat ad hoc and the performance of a given strategy may vary from one application to another. The main aim of a general method is that it gives fast convergence and robust stable estimates over a wide

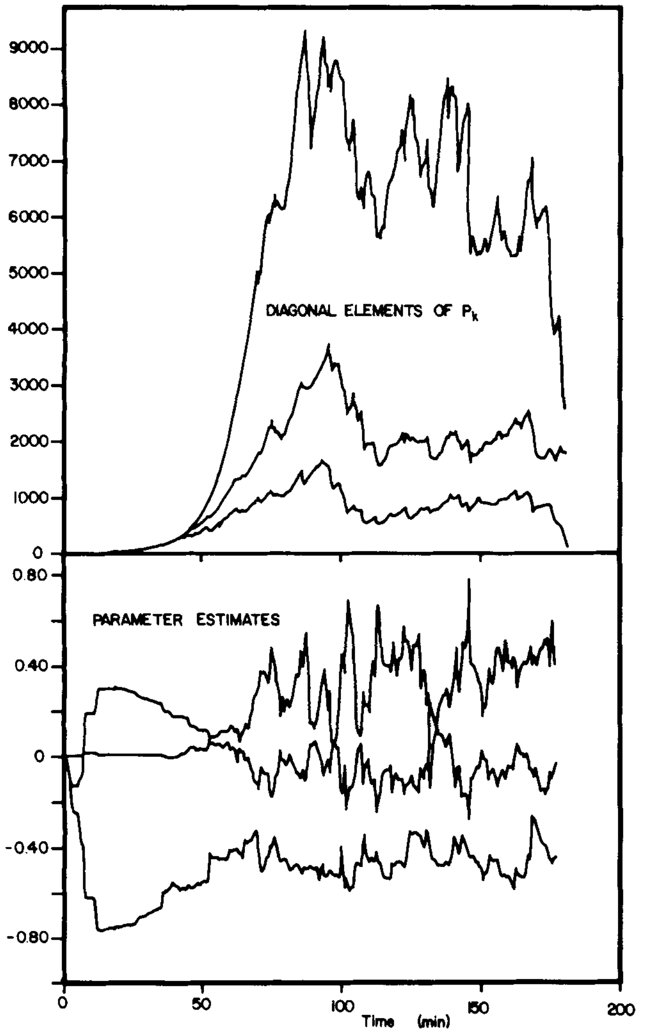

Figure 1. Blow-up of P_k and parameter estimates due to constant forgetting factor, $\lambda = 0.99$; sampling time = 6 s.

range of operating conditions. Furthermore, we desire that the scheme have design parameters that are easy to choose and few in number.

The minimum-variance type estimator has been reported to give about as fast a convergence as can be expected from a recursive estimator (Kumar and Moore, 1982; Fogel, 1982).

Weighted minimum variance estimates of θ can be obtained recursively by the following algorithm:

Algorithm 1.

Given

$$\{P_o > 0, T \geq 1, \hat{\theta}_o\}$$
$$\{r_k > 0, 0 < \lambda_k \leq 1; k \geq 1\}$$

set $k = 1$

Compute: 1. $P_k^{-1} = \lambda_k P_{k-1}^{-1} + \phi_{k-T}\phi'_{k-T}r_k^{-1}$

2. $\hat{\theta}_k = \hat{\theta}_{k-1} + P_k\phi_{k-T}r_k^{-1}[y_k - \phi'_{k-T}\hat{\theta}_{k-1}]$

Set $k = k + 1$ then go to step 1.

$\hat{\theta}_k$ in the algorithm is the conditional estimate of θ based on the available information at time k, P_k is the corresponding covariance matrix, r_k the variance of measurement y_k, and λ_k is a forgetting factor designed to increase the variance of measurements as they age. If the values of (u_k, y_k) are missing for $k < 0$ it suffices to take the missing values as zero; the algorithm then holds $\hat{\theta}_k = \hat{\theta}_o$ until data become available. (The algorithm can also be reformulated to avoid the matrix inversion called for in step 2.)

We can see that for $\lambda_k = 1$, P_k^{-1} is continually increasing which results in a decreasing *sensitivity* to new data in Step 2 and the algorithm eventually loses its adaptability. A fixed forgetting factor less than unity prevents this undesirable behavior, but may cause other problems, such as decreasing P_k^{-1} when there is little information in the data; this can lead to a drift in the estimates or to instability. A typical experimental result is illustrated in Figure 1. Fortescue et al. (1981) suggest that λ_k be selected in order to keep a measure of the information content of the estimation constant. A suitable measure is the "nominal memory length," N_k, which is a weighted sum of squares of a posteriori errors, that is, the weighted residuals (obtained using the current estimate of parameters) from all past data scaled by their respective measurement variance,

$$N_k \equiv \sum_{i=1}^{k} \sigma_{i/k} r_i^{-1}(y_i - \phi'_{i-T}\hat{\theta}_k)^2 \tag{17}$$

where

$$\sigma_{k/k} = 1;\ \sigma_{i/k} = \lambda_k \sigma_{i/k-1};\ i = 1,2, \ldots k - 1$$

It can be shown that, after some considerable manipulation, Eq. 17 can be updated recursively (Ydstie, 1982) to yield

$$N_k = \lambda_k N_{k-1} + (1 - \phi'_{k-T}P_k\phi_{k-T}r_k^{-1})r_k^{-1}e_k^2$$

or, more conveniently, in terms of the previous value of the covariance,

$$N_k = \lambda_k N_{k-1} + \lambda_k e_k^2/(r_k\lambda_k + \phi'_{k-T}P_{k-1}\phi_{k-T}) \tag{18}$$

where

$$e_k \equiv y_k - \phi'_{k-T}\hat{\theta}_{k-1}$$

In the proposed algorithm, λ_k is chosen to maintain a constant value for N_k. Thus at each step enough past information is discounted to insure that each estimation of parameters leads to the same weighted sum of residuals and hence, in some sense, is based on the same amount of information. Thus, setting $N_k = N_{k-1} = \ldots = N_o > 0$ in Eq. 18, we obtain

$$\lambda_k = (n_k + \sqrt{n_k^2 + 4w_k})/2 \tag{19}$$

where

$$w_k \equiv \phi'_{k-T}P_{k-1}\phi_{k-T}r_k^{-1};\ n_k \equiv 1 - w_k - e_k^2/(r_kN_o)$$

It is easy to demonstrate that λ_k calculated by Eq. 19 satisfies $0 < \lambda_k \leq 1$ for all k and that λ_k is close to unity if any of the following is true: the process is not excited; the parameters $\hat{\theta}_k$ are close to their correct values; the uncertainty in $\hat{\theta}_{k-1}$ is large (w_k large); the parameter N_o is large. On the other hand, values of the λ_k significantly less than unity occur when the error e_k is large but P_{k-1} has a lower magnitude, leading to small w_k. This, in turn, causes an increase in P_k leading to a retuning of parameters.

In addition, Ydstie (1982) and Ydstie and Sargent (1982) have shown that the estimator, when coupled with a minimum variance regulator, possesses desirable convergence and stability properties in deterministic and certain stochastic cases.

For practical implementation of the estimator on real plants, one would not generally have knowledge of the noise variance r_k, and how it might vary with time. It can be shown, however, that if the noise variance is assumed to be constant and the initial data are ignored ($P_o \rightarrow \infty$), the estimation procedure reduces to a recursive least-squares algorithm in which the "covariance" matrix P is scaled by the (unknown) noise variance r. A suitable information measure in this case is given by a dimensional form of the memory length,

$$\Sigma'_k = r \cdot N_k = \sum_{i=1}^{k} \sigma_{i/k}(y_i - \phi'_{i-T}\hat{\theta}_k)^2 \qquad (20)$$

This least-squares formulation of the estimation algorithm was used in all of our simulation and experimental pilot plant studies.

Closed-Loop Algorithm: Estimation Plus Control

The estimation procedure described above can be used in conjunction with either a minimum variance or an extended horizon control strategy to provide an adaptive self-tuning regulator. Of course, the various proofs of convergence and stability of the resulting algorithms are valid only for linear, time-invariant processes and not for real plants.

Furthermore, identifiability of the closed-loop system may be lost when measurements go out of range or when valves saturate. In theory, therefore, it is necessary to stop estimation when this occurs. Another consideration is that the system is not generally closed-loop identifiable when it is operated with a constant feedback control law (Åström and Wittenmark, 1973) and those authors suggest fixing one of the β parameters. Ljung (1977), however, shows that the only fixed point solutions of the algorithm are those that yield optimal control. In addition, simulation and plant experiments by many workers have shown that the above considerations may not be important in the context of adaptive control of real process plants.

The major limitations of the minimum variance control strategy are that the delay time must be known and be constant, and that the system must have a stable inverse. Both these restrictions can be relaxed when minimum variance control is replaced by the extended horizon control strategy. Only an upper bound for the system delay time is required and a wide class of nonminimum phase systems can be stabilized by using the detuning property of the time horizon (Ydstie, 1982; Ydstie and Liu, 1984; Hiram and Kershenbaum, 1983).

The complete algorithm can be obtained in a convenient form from the above equations after some algebraic manipulations to yield:

Algorithm 2. Extended Horizon Adaptive Control.

Given: $\{P_o > 0,\ T \geq$ maximum delay time, $\hat{\theta}_o, \Sigma_o > 0,\ n \geq 1\}$

$$\{u_k, y_k;\ k \geq 1\}$$

Set $k = 1$

Compute:
1. $e_k = y_k - \phi'_{k-T}\hat{\theta}_{k-1}$
2. $w_k = \phi'_{k-T}P_{k-1}\phi_{k-T};\quad n_k = 1 - w_k - e_k^2/\Sigma_o$
3. $\lambda_k = (n_k + \sqrt{n_k^2 + 4w_k})/2$
4. $K_k = P_{k-1}\phi_{k-T}/(\lambda_k + w_k)$
5. $\hat{\theta}_k = \hat{\theta}_{k-1} + K_k e_k$
6. $P_k = [P_{k-1} - K_k K'_k(\lambda_k + w_k)]/\lambda_k$
7. $u_k = (y^*_{k+T} - h_{k+T}) \Big/ \sum_{j=1}^{T} \beta_j$

Set $k = k + 1$, go to step 1.

Comments

(i) In step 7, h_{k+T} is given by Eq. 8 above and both h_{k+T} and $\Sigma\beta_j$ are calculated using the current estimate of the parameters, $\hat{\theta}_k$.

(ii) Also in step 7, we have used the control strategy 1, Eq. 9. This can be replaced by other strategies or the scheme can use the incremental model and control law, Eqs. 14–16.

(iii) Wide experience from simulation of plants that exhibit changes in dynamics indicates that it is preferable to use factored updates of P_k rather than updates of P_k directly as indicated by step 6. In most of our work, updating of $P^{1/2}$ was carried out with a modification of the algorithm by Peterka (1975). Biermann (1977) discusses this in detail.

(iv) On computers with small word lengths, calculation of λ_k via steps 2 and 3 may be numerically unsuitable. An alternative method is to solve Eq. 18 for λ_k by successive substitution with $N_k = N_{k-1} = N_o$ and an initial guess of $\lambda_k = 1$. Adequate convergence is achieved within five iterations and the time taken is comparable with that required to extract the square root in step 3.

(v) The algorithm contains three main adjustable parameters: n, T, and Σ_o; the sensitivity of the algorithm to the choice of these parameters is discussed below.

(vi) The algorithm has been formulated using the least-squares estimation of parameters for the reasons given earlier. No knowledge of the noise characteristics is required, although some idea of r is helpful in choosing Σ_o (see Eq. 20) which can be interpreted as an effective memory length times the expected variance of the predictions.

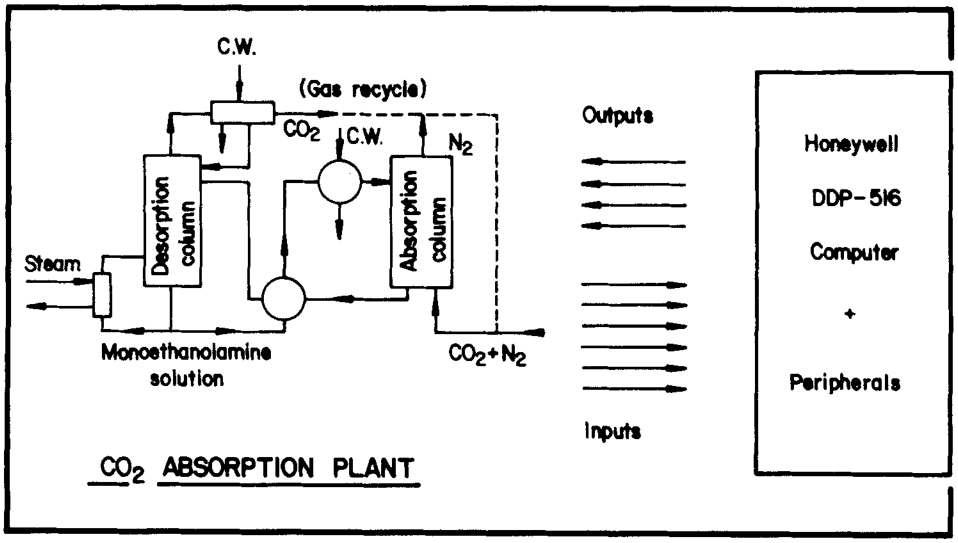

Figure 2. Schematic diagram of pilot plant and computer system.

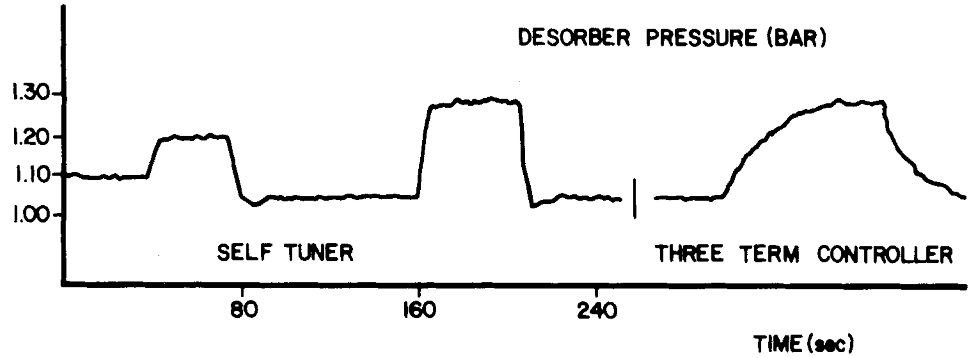

Figure 3. Comparison of minimum variance self-tuning regulator with the best PID controller for a step change in set point.

EXPERIMENTAL STUDIES

The plant and computer system as used in this research program is illustrated in Figure 2. (The computer has now been replaced by an IBM-4341 system.) The carbon dioxide absorber-desorber unit consists of two columns (9 m high × 0.25 m diam) for separation of CO_2 from nitrogen using an ethanolamine solution as absorbent. The CO_2 is continuously absorbed in the absorption column [at 2–4 bar (200–400 kPa)] and is stripped in the desorption column, usually at atmospheric pressure. The loops that have been used to study some of the characteristics of the algorithms are flow, pressure, level, and composition loops. The aim is to find a controller that will outperform a well-tuned PID controller for a wide range of operating conditions.

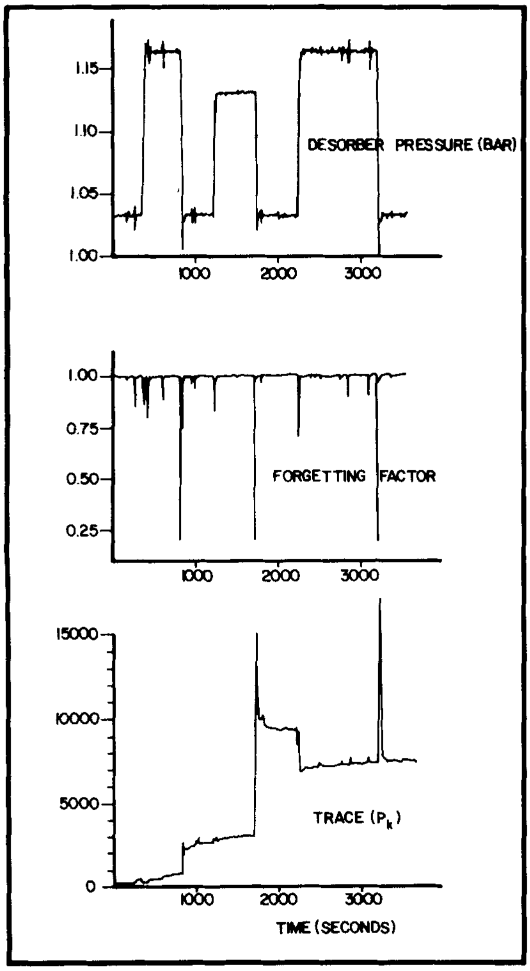

Figure 4. Algorithm behavior (minimum variance control) during a sequence of set point changes ($\Sigma_o = 0.005$, $N_O \sim 10$).

Of special interest is the pressure control of the desorption column. This process is nonlinear and is highly interactive with the pressure in the absorption column when CO_2 is fed back through the compressor unit as indicated by the dotted line. Control is achieved via a control valve in a recycle line around the compressor. Off-line identification studies using correlation techniques and pseudorandom excitation show that a third-order model sampled at 4–6 s gives a reasonable description of the pressure dynamics (Tranmer, 1983). An identified discrete-time normalized model (sampling time = 6 s) is

$$y_k - 1.044y_{k-1} + 0.097y_{k-2} + 0.073y_{k-3} = -0.026u_{k-1} - 0.019u_{k-2} + 0.177u_{k-3},$$

which has nonminimum phase characteristics.

In initializing the algorithm, estimation was started first and after 3–6 estimator steps, the control loop was closed. (Later results have shown, however, that performnace will be as good or better if the estimator and controller are started simultaneously—Hiram and Kershenbaum, 1983.) The self-tuner applied to this system converged quickly due to the deterministic nature of the problem and easily outperformed any PI regulator tuned to perform well for a wide range of process conditions (see Figure 3). Similar results were achieved for flow and level control (Ydstie, 1982; Ahmad, 1981); level control was achieved by manipulating the flow rate out of the base of the columns. The self-tuner performed at least as well as a well-tuned PID regulator and had the additional advantage of a high degree of robustness and ability to adapt to changing parameters.

As pointed out by Ydstie and Sargent (1982), adaptability is closely related to the size of P_k. Simulation and plant experiments, as well as analysis, show that the steady-state value of P_k is inversely related to choice of Σ_o. A large Σ_o will give a small P_k and vice versa. However, changes in the process dynamics lead to large variations in P_k and, as indicated in Figure 4, P_k can show a spiky behavior when there are abrupt changes and when Σ_o is chosen small.

The robustness with respect to design parameters of the minimum variance and extended horizon adaptive controllers has been investigated in an extensive program (Ahmad, 1981; Ydstie, 1982) involving responses to load and set point changes as well as to stochastic disturbances. The improved robustness of the extended horizon algorithm is clearly shown in Figures 5–8.

The results of a typical sequence of experiments using the minimum variance self-tuning regulator ($T = D = 1$) with variable forgetting factor is illustrated in Figure 5; three of the adjustable parameters—the gain β_1 (here β_1 was fixed; in all other runs it was estimated), the model order, and the sample time—were changed in turn. Some measure of robustness to choice of parameters was demonstrated and it was usually possible to find parameters that would yield stable control for the pressure, level, and flow loops.

Figure 6 illustrates the best achievable performance for pressure control with minimum variance regulation. With this control law, some overshoot and steady-state ripple occurred with every choice of parameters because of the high performance objective. Furthermore, for both pressure and flow control a unique value of the delay time ($D = 1$) had to be selected in order to achieve satisfactory performance with a minimum variance control law;

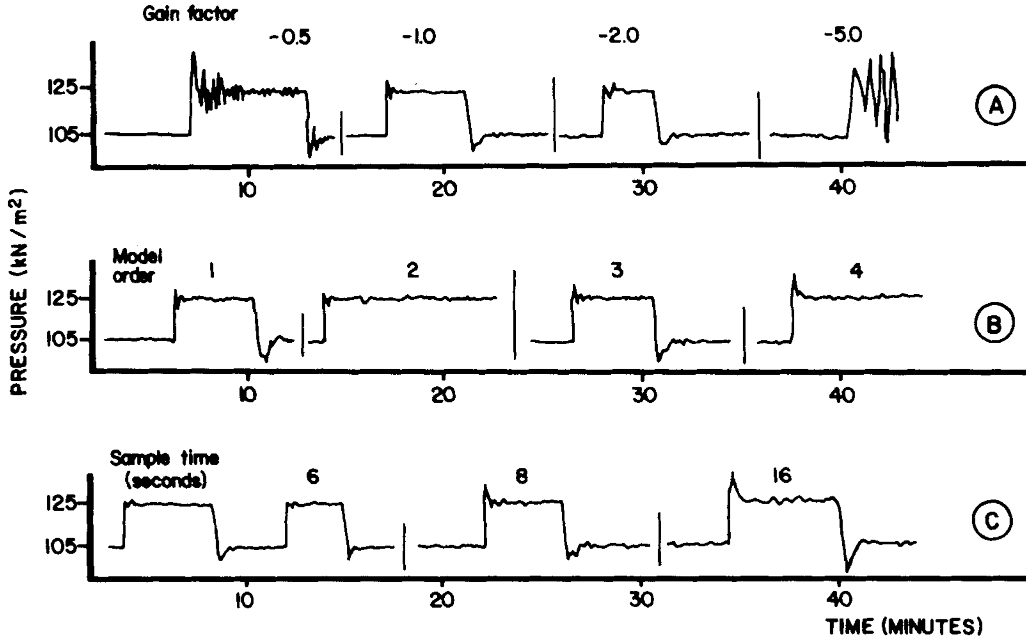

Figure 5. Sensitivity of minimum-variance self-tuning regulator to choice of parameters (pressure control) (a) estimated gain; (b) model order; (c) sample time.

for all other choices of delay time ($D > 1$), convergence was difficult or impossible to attain.

The extended horizon controller, on the other hand, appears to overcome these problems without significantly reducing dynamic performance for $T = 2$ (Figure 7). A time horizon $T = 3$ gives similar performance to the PI regulator shown in Figure 3 but one which will remain equally good upon change of operating conditions. Larger values of the time horizon also gave satisfactory albeit somewhat more sluggish, performance; the dynamic performance can be significantly improved, however, by decoupling the sampling rates for estimation and control and implementing the control (step 7 of algorithm 2) more frequently than the estimation.

Figure 8a illustrates the effect of choice of delay time with the minimum variance control strategy for level control; in this case the choice of D is not so crucial as it was for pressure and flow control. The intersample ripple can be reduced by choosing a longer sample interval at the cost of more sluggish response. Alternatively, as shown in Figure 8b, an extended horizon controller with the same sample time but a relatively long time horizon can be used; control remains excellent despite variations in choice of T.

DISCUSSION OF RESULTS

The type of experiments illustrated in the previous section were carried out over a wide range of conditions and choice of design parameters. In the following we summarize the results which indicate that several of the adaptive strategies show considerable robustness and flexibility. The results are discussed in terms of the sensitivity of the various algorithms to choice of design parameters or initial estimates. Most discussion applies to the extended horizon controller (Algorithm 2) which had the most promising behavior over the widest range of operating conditions. Some discussion of further developments is also included.

Initial Estimates. A range of different initial estimates was tried and the algorithms were found to be totally insensitive to choice of these. An arbitrary guess of $\hat{\theta}_o$ was used in most cases and found satisfactory provided update of control action was suspended until the parameters adjusted to nonzero values (to prevent division by zero or rapid oscillation in signs of the parameters). The initial guess for the covariance matrix was, in general, chosen to be very large ($10^6 \times I$, or larger) and this was found to function well both in simulation and plant experiments. Small initial covariance matrices lead to slow initial convergence.

Model Order. The adaptive algorithms were not sensitive to choice of model order. Due to the simple nature of the problem, acceptable performance was obtained even with first-order models. Third-order models, in general, gave better performance, but no further improvement was seen with higher order models.

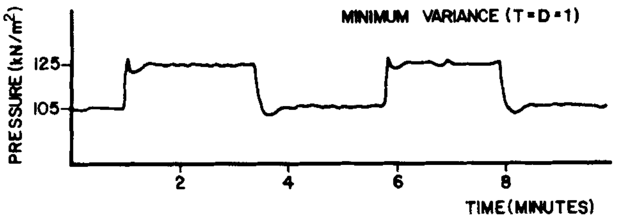

Figure 6. Performance of minimum variance self-tuning regulator (pressure control).

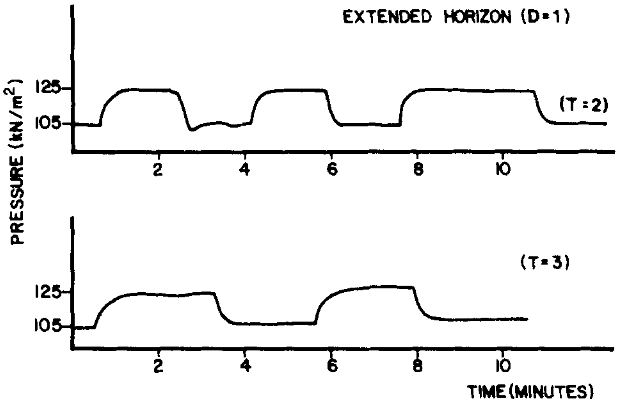

Figure 7. Sensitivity of extended horizon self-tuning regulator to choice of time horizon (pressure control).

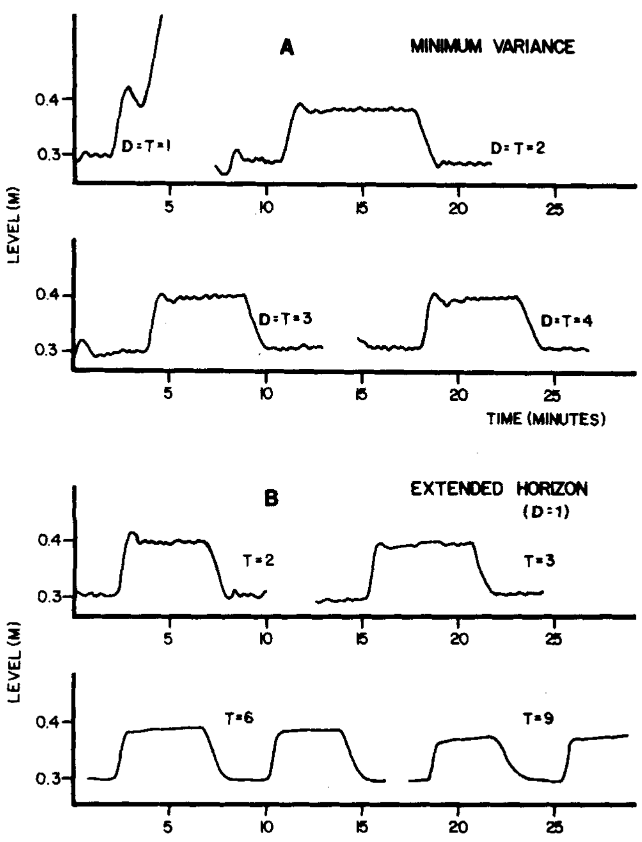

Figure 8. Sensitivity of self-tuning regulators to choice of delay time/time horizon (level control): (a) minimum variance control; (b) extended horizon control.

Sample Time. Various sample rates were tried; 4–6 s gave the best performance for pressure and flow control and 8–12 s for the level control. The self-tuner, however, gave stable control for a wide range of sample periods (2–32 s for the pressure loop), but intersample drifts and oscillations, as well as overshoot during set point changes, increased with increasing sample time. Extended horizon control showed even less sensitivity.

Choice of Σ_o. A large Σ_o gives small steady-state covariance matrices and slow adaptation, whereas a small Σ_o gives a more responsive system and rapid adaptation at the cost of having larger parameter uncertainty and steady-state process variance. $\Sigma_o = 0.25$ was found to work satisfactorily on the processes considered here; this is 100–1,000 times the estimated variance of the measurements. The sensitivity to choice of Σ_o was, however, found to be small and numerous experiments and computer simulations confirm that the estimator with variable forgetting factor is able to adapt both to slow and sudden changes in process dynamics for a wide range of Σ_o (several orders of magnitude).

Choice of β_1. This is only relevant with a minimum variance control law and when the strategy of Åström and Wittenmark (1973) is used to fix one parameter to avoid "free" drift of parameters in the approximation space. It was found that $\beta_1 = -1.0$ gave satisfactory performance for the normalized system. The range of β_1 for which we obtained stable control was, in some cases, narrow (a factor of 10 or less) but difficulties in guessing a reasonable value of β_1 by trial and error were only slight. However, a much better approach is to estimate β_1 along with the rest of parameters of θ as suggested in Algorithm 2; no difficulties were ever experienced when this was done and much time was saved in starting up the algorithm.

Delay Time: The choice of the right delay time was found to be important for the minimum variance control algorithm with fixed β_1. The algorithm that estimates all parameters was found to be somewhat less sensitive to choice of this as long as it was chosen at least as large as the transportation delay in the system. The extended horizon controller provided a much better solution by not requiring an estimate of the delay time at all. A number of leading coefficients of $B(z^{-1})$ were estimated (on-line) to be close to zero depending upon the actual delay time in the system; no intervention by the user was required.

Time Horizon. A long time horizon requires the on-line estimation of more parameters and gives slower response, but leads to more robust control. Time horizons between 2 and 8 sample intervals gave satisfactory performance for the loops listed here; in most cases, the actual delay time in the systems studied was less than 2 s.

Incremental Algorithm. The incremental algorithm gave both good transient behavior and zero steady-state error and is to be preferred when set point following is important.

Numerical Considerations. Because of the accumulation of round-off errors, direct implementation of Algorithm 2 (specifically step 6) is not satisfactory unless nonlinearities are mild and the period of operation is limited. Instabilities could be induced after only 200–300 steps when a time-varying, nonstationary noise system was used. However, the use of the square root factored algorithm for covariance update (Peterka, 1975) insured that P_k always stayed nonnegative and avoided the numerical problems in all cases studied; this form of the update is recommended for other users.

Noise Characteristics. The theoretical discussion in this paper has centered around the adaptive control of a white noise system. Measurements taken from several loops of the CO_2 pilot plant suggest that this gives a realistic noise model for most, but not all, of the processes considered here (Fortescue, 1977; Albrecht et al., 1980). In current control literature, the discussion has to a large extent concerned systems that are perturbed by colored noise v_k' generated by a linear filter such that $v_k' = C(z^{-1})v_k$ where v_k is as before and $C(z^{-1})$ is a monic and stable polynomial. The estimation problem is now nonlinear and the AML method of Solo (1979) or similar techniques can be used to estimate $C(z^{-1})$. Ljung (1977), however, shows that the self-tuner without estimating $C(z^{-1})$ converges to the optimal law and the extra effort involved hardly seems merited (Fogel, 1982). Our experience tends to support that conclusion.

Multivariable Cases. The algorithms here are readily extendable to cover multivariable problems (Goodwin et al., 1980; Ydstie, 1982; Ydstie and Liu (1984) and certain classes of nonlinear problems (Svornos et al., 1981). The number of parameters to be estimated, however, increases rapidly and it is important, especially in the multivariable case, to take account of the natural block structure of the canonical models so that parameter redundancy is avoided. This problem will be discussed in future publications.

ACKNOWLEDGMENT

The authors wish to thank British Petroleum Ltd., who generously provided a research studentship for one of the authors (B.E.Y.) to carry out this research. We would like to acknowledge the work of Mr. A. Ahmad, who carried out many of the pilot plant experiments, and of Dr. Y. Hiram, who assisted with the simulations.

NOTATION

A,B,C	= polynomials in z^{-1}
D	= estimated model delay time
d	= bias parameter
e_k	= innovation sequence

$E(\cdot),E(\cdot\|Y_k)$	= expected value and conditional expectation
F,G	= polynomials in z^{-1}
h_k	= predicted output in extended horizon control law
K_k	= estimator gain
N_k	= nominal memory length
n	= model order
n_k	= scalar defined by Eq. 19
P_k	= covariance matrix
q_k,r_k	= noise variances
T	= time horizon
u_k	= input sequence
v_k	= noise sequence
w_k	= scalar defined by Eq. 19
y_k	= output sequence
y_k^*	= reference output sequence
Y_k	= collection of all available information up to time k
z^{-1}	= backward shift operator

Greek Letters

α,β,δ	= parameters in prediction model
$\Delta \equiv 1 - z^{-1}$	= incremental operator
ϵ_k	= noise sequence
θ	= vector of parameters
λ_k	= forgetting factor
$\sigma_{i/k}$	= product of forgetting factors
Σ_k	= weighted sum of residuals

Superscripts

$\bar{}$	= refers to incremental model
$\hat{}$	= estimated value

Subscripts

k	= kth time interval

LITERATURE CITED

Ahmad, A., "Application and Comparison of Different Self-Tuning Regulator Algorithms on a Pilot Plant," M.Sc. Thesis, Imperial College, London (1981).

Albrecht, J. J., L. S. Kershenbaum, and D. L. Pyle, "Identification and Linear Multivariable Control in an Absorption-Desorption Pilot Plant," *AIChE J.*, **26**, 496 (1980).

Åström, K. J., et al., "Theory and Application of Self-Tuning Regulator," *Automatica*, **13**, 457 (1977).

Åström, K. J., and B. Wittenmark, "On Self-Tuning Regulators," *Automatica*, **9**, 185 (1973).

———, "Self-Tuning Controllers Based on Pole-Zero Placement," *Proc. IEE, Pt. D*, **127**, 120 (1980).

Bierman, G. J., *Factorization Methods for Discrete Sequential Estimation*, Academic Press, New York (1977).

Clarke, D. W., and P. J. Gawthrop, "Self-Tuning Controller," *Proc. IEE*, **122**, 929 (1975).

———, "Self-Tuning Control," *Proc. IEE*, **126**, 633 (1979).

———, "Implementation and Application of Microprocessor-Based Self-Tuners," *Automatica*, **17**, 233 (1981).

Clarke, D. W., A. J. F. Hodgson, and P. S. Tuffs, "Offset Problem and k-Incremental Predictors in Self-Tuning Control," *Proc. IEE, Pt. D*, **130**, 217 (1983).

Cutler, C. R., and B,. L. Ramaker, "Dynamic Matrix Control—A Computer Control Algorithm," *Proc. JACC*, WP5-B, San Francisco (1980).

Doss, J. E., et al., "New Directions for Process Control in the Eighties," AIChE Diamond Jubilee Meet., Washington, DC (1983).

Dumont, P., "Self-Tuning Control of a Chip Refiner Motor Load," *Automatica*, **18**, 307 (1982).

Dumont, G. A., and P. R. Bélanger, "Successful Industrial Application of Advanced Control Theory to a Chemical Process," *IEEE Control System*, **1**, 1 (1981).

Fjeld, M., B. Foss, and S. Aam, "Regulation of River Water Level by Adaptive Control of a Reservoir Trap," *Proc. ACC*, San Francisco (1983).

Fjeld, M., et al., "Application of a Self-Tuning Regulator and a Continuous Performance Monitor to Paper Machine Control," AccuRay Corp. Report (1981).

Fogel, E. "On Self-Tuning Regulators and the Least-Squares Procedure," *CSDL-P-1420*, Charles Stark Draper Laboratory, Cambridge, MA (1982).

Fortescue, T. R., "Work on Åström's Self-Tuning Regulator," Dept. of Chem. Eng. Report, Imperial College, London (1977).

Fortescue, T. R., L. S. Kershenbaum, and B. E. Ydstie, "Implementation of Self-Tuning Regulators with Variable Forgetting Factors," *Automatica*, **17**, 831 (1981).

Goodwin, G. C., and L. Dugard, "Stochastic Adaptive Control with Known and Unknown Interactor Matrices," IFAC Symp. Adaptive Control, San Francisco (1983).

Goodwin, G. C., P. J. Ramadge, and P. E. Caines, "Discrete Time Multivariable Adaptive Control," *IEEE T-AC-25*, 449 (1980).

Goodwin, G. C., and K. S. Sin, "Adaptive Control of Nonminimum Phase Systems," *IEEE T-AC-26*, 478 (1981).

Hiram, Y., and L. S. Kershenbaum, "Simulation Studies of Extended Horizon Self-Tuning Regulators," Dept. of Chem. Eng. Report, Imperial College, London (1983).

Kumar, R., and J. B. Moore, "Convergence of Adaptive Minimum Variance Algorithms via Weighting Coefficient Selection," *IEEE-T-AC-27*, 146 (1982).

Lee, K. S., and W. K. Lee, "Extended Discrete Time Multivariable Adaptive Control Using Long Term Predictor," *Int. J. Control*, **38**, 495 (1983).

Ljung, L., "On Positive Real Transfer Functions and the Convergence of Some Recursive Schemes," *IEEE-T-AC-22*, 539 (1977).

Mehra, R. K., et al., "Model Algorithmic Control (MAC): Review and Recent Developments," Sea Island, Georgia, *Chemical Process Control 2*, Edgar & Seborg, Eds., 297 (1982).

Mosca, E., and G. Zappa, "Removal of Positive Realness Condition in Minimum Variance Adaptive Regulators by Multistep Horizons," *IEEE Trans*, **AC-29**, 844 (1984).

Peterka, V., "Adaptive Digital Regulation of Noisy Systems," 2nd IFAC Symp. on Identification and Process Parameter Estimation, Prague (1970).

———, "A Square Root Filter for Real-Time Multivariable Regression," *Kybernetika*, **11**, 53 (1975).

Richalet, J., et al., "Model Predictive Heuristic Control: Applications to Industrial Processes," *Automatica*, **14**, 413 (1978).

Sastry, V. E., D. E. Seborg, and R. K. Wood, "An Application of a Self-Tuning Regulator to a Binary Distillation Column," *Automatica*, **13**, 417 (1977).

Seborg D. E., S. L. Shah, and T. F. Edgar, "Adaptive Control Strategies for Process Control: A Survey," AIChE Diamond Jubilee Meet., Washington, DC (1983).

Solo, V., "The Convergence of AML," *IEEE T-AC-24*, 958 (1979).

Svornos, S., G. Stephanopoulos, and R. Aris, "On Bilinear Estimation and Control," *Int. J. Control*. **34**, 651 (1981).

Tranmer, J., "Identification of Multivariable Discrete Time Transfer Function Models for a CO_2 Absorption Plant," Dept. of Chem. Eng. Report, Imperial College, London (1983).

Tung, L. S., "Sequential Predictive Control of Industrial Processes," *Proc. ACC*, San Francisco (1983).

Wellstead, P. E., et al., "Self-Tuning Pole-Zero Assignment Regulators," *Int. J. Control*, **30**, 1 (1979).

Wittenmark, B., "A Self-Tuning Regulator," Lund Univ. Report, Lund, Sweden (1973).

Ydstie, B. E., "Robust Adaptive Control of Chemical Processes," Ph.D. Thesis, Univ. of London (1982).

———, "Extended Horizon Adaptive Control," 9th Triennial IFAC World Cong., Budapest (1984).

Ydstie, B. E., and L. K. Liu, "Extended Horizon Prediction and Adaptive Multivariable Control," *Proc. ACC*, San Diego (1984).

Ydstie, B. E., and R. W. H. Sargent, "Deterministic Convergence of an Adaptive Regulator with Variable Weighting of Past Data," 6th IFAC Symp. on Identification and System Parameter Estimation, Washington, DC (1982).

Manuscript received July 7, 1983; revision received Jan. 3, and accepted Jan. 8, 1985.