1 of 1

Conf-940307--8

LA-UR-94-0171

Title: PROCESS FAULT DETECTION AND NONLINEAR TIME SERIES ANALYSIS FOR ANOMALY DETECTION IN SAFEGUARDS

Author(s): T. L. Burr and L. E. Wangen

Submitted to: IAEA Symposium, Vienna, Austria, March 14-18, 1994

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Los Alamos
NATIONAL LABORATORY

PROCESS FAULT DETECTION AND NONLINEAR TIME SERIES ANALYSIS FOR ANOMALY DETECTION IN SAFEGUARDS

T. L. Burr, M. F. Mullen, L. E. Wangen
Los Alamos National Laboratory

ABSTRACT

In this paper we discuss two advanced techniques, process fault detection and nonlinear time series analysis, and apply them to the analysis of vector-valued and single-valued time-series data. We investigate model-based process fault detection methods for analyzing simulated, multivariate, time-series data from a three-tank system. The model-predictions are compared with simulated measurements of the same variables to form residual vectors that are tested for the presence of faults (possible diversions in safeguards terminology). We evaluate two methods, testing all individual residuals with a univariate *z*-score and testing all variables simultaneously with the Mahalanobis distance, for their ability to detect loss of material from two different leak scenarios from the three-tank system: a leak without and with replacement of the lost volume. Nonlinear time-series analysis tools were compared with the linear methods popularized by Box and Jenkins. We compare prediction results using three nonlinear and two linear modeling methods on each of six simulated time series: two nonlinear and four linear. The nonlinear methods performed better at predicting the nonlinear time series and did as well as the linear methods at predicting the linear values.

1. INTRODUCTION

For process control and other reasons, there is increasing automation of chemical processing plants, including spent-nuclear-fuel reprocessing plants. Consequently there will be more data *potentially* available for safeguards in future reprocessing. These data will consist of control data and physical and chemical measurements of process inputs and outputs during plant operations. Not only will more variables be monitored, but data collection will be more frequent than in the past. These data could assist the safeguards function if appropriate data analysis methods can be identified. We have investigated two different approaches for the analysis and interpretation of such data: 1) process fault detection applied to the monitoring of multivariate time-series data and 2) application of nonlinear methods to the analysis of univariate time-series data. The latter differs from the present linear methods used for time series analysis in safeguards.

The first method, **process fault detection and diagnosis**, monitors a vector-valued time series, such as the amount of nitric acid, plutonium, and uranium in various tanks and other vessels over time. The second method, **nonlinear time-series analysis**, monitors a scalar-valued time series, such as the amount of plutonium over time, but allows the serial dependence to have an arbitrary functional form. Present methods used to analyze time-series data for materials control and accountancy, such as material unaccounted for (MUF), assume that the functional form of the time series is linear.

The goal of **process fault detection** and diagnosis is to develop improved methods for detecting, isolating, and identifying deviations from nominal or desired process operating conditions [1]. Process fault detection and diagnosis involves comparing data from process measurements with redundant information to detect and identify faults so that appropriate action can be taken. The redundant information can be either from other process measurements or from process models. The concept is illustrated in Fig. 1 for a case in which a process model is used to provide the redundant information. The process model is developed for "normal" operating conditions, and it uses process knowledge and process inputs to make predictions about the expected state of selected output variables. These predictions are compared with measured values of the same output variables to form residuals (the differences between the measured and model-predicted values) that are tested to determine the presence or absence of a fault at a desired degree of confidence. Because of the unavoidable presence of both modeling and measurement errors, non-zero residuals are expected so criteria are needed for deciding whether a fault has occurred. Mass balance relations in the form of MUF are an example of a process model involving simple consistency relationships. In this report we focus on the analysis of multivariate residuals representing a single point in time. In another report [2] we address the analysis of cumulative residuals from successive multivariate residuals.

We compared **nonlinear time-series analysis tools** [3,4] with the linear methods popularized by Box and Jenkins [5] for analyzing time series. Box-Jenkins analysis assumes that the expected value of any observation is a linear function of some subset of all previous observations and errors. Our approach relies on using historical data to estimate the functional dependence of the present observation on some subset of past observations. We then estimate the expected value of the next observation in the time series using a nonparametric procedure (no distribution assumptions) based on past observations. The next value of the time series is predicted, then compared with the observed value. If the resulting residual is large, we suspect that an anomaly has occurred. We illustrate the idea in Fig. 2. Shown in Fig. 2a is a time series plot of the first 100 values of a nonlinear time series. In Fig. 2b, is our conditional expected value estimate (solid line) on a scatter plot of the present versus the past values from the same time series. Note that the nonlinear dependence of the present value on the past value is not readily apparent from Fig. 2a but is apparent in Fig. 2b and that our estimated conditional mean can be viewed as being a scatter plot smoother in the lag one case. This estimated conditional mean can be used together with the estimated prediction error to assist with detecting anomalies.

2. METHODS AND RESULTS: PROCESS FAULT DETECTION

We applied two multivariate fault detection techniques to simulated data from a three-tank system (Fig. 3) containing nitric acid, plutonium, and uranium. The dynamics are described by a system of coupled differential equations based on total mass balances and on individual mass balances for each chemical species for each tank (Eq. 1).

$$[\text{Time rate of change of mass}] = [\text{Mass in}] - [\text{Mass out}] \qquad (1)$$

For given input flows, initial tank volumes, and initial concentrations of nitric acid, plutonium, and uranium, the differential equations are solved to give the outputs, i.e., the volumes and concentrations in the tanks at various times. When density is a linear function of concentration, the equations for Tank 1 are

$$
\begin{aligned}
dV_1/dt &= F_{11} + F_{12} - F_{21} \\
dH_1/dt &= [(H^0_{11}\, F_{11} + H^0_{12}\, F_{12}) - (F_{11} + F_{12})\, H_1] \,/\, V_1 \\
dPu_1/dt &= [(Pu^0_{11}\, F_{11} + Pu^0_{12}\, F_{12}) - (F_{11} + F_{12})\, Pu_1] \,/\, V_1 \\
dU_1/dt &= [(U^0_{11}\, F_{11} + U^0_{12}\, F_{12}) - (F_{11} + F_{12})\, U_1] \,/\, V_1
\end{aligned}
$$

with analogous equations for Tanks 2 and 3. The superscript zeros are tank input concentrations for nitric acid (H), plutonium (Pu), and uranium (U), and the other symbols are defined in Fig. 3. The density of each tank solution is determined from empirical relationships between density and concentrations of nitric acid, plutonium and uranium [6,7]. The system of equations is solved by the Euler method. The volumes, densities, and concentrations of plutonium, uranium, and nitric acid are the model predictions that are compared to measured values to give residuals, which are tested for faults. Simulated measured values are obtained by adding the following relative standard deviations to the known true values: flow rate–0.05, tank volume–0.002, density–0.002, and nitric acid, plutonium, and uranium concentrations—0.01, 0.002, and 0.004—respectively.

The residuals were evaluated by two different multivariate techniques. The first technique monitors each individual element of the residual vector separately and is a natural extension of the commonly used univariate approach. The univariate test statistic for variable p of vector-valued residual i, r_{ip}, that is expected to be zero, $\mathbf{H}_0$: $E(r_{ip}) = 0$ versus the alternative hypothesis $\mathbf{H}_1$: $E(r_{ip}) \neq 0$, is

$$z_{ip}(\alpha) = \frac{r_{ip}}{\sigma/\sqrt{n}} \,, \qquad (2)$$

assuming the standard deviation σ is known. Here, n is the number of samples used to calculate $\mathbf{r}$ ($n = 1$ for this work) and E denotes expected value. The critical values to which these test statistics are to be compared come from N(0,1). The user specifies what significance level (α value) will be used to signal a fault depending on the number of false alarms to be tolerated. For uncorrelated, multivariate normal distributions, if we wish to maintain the same overall significance level for detecting a fault, tests for individual residuals use the Bonferroni method [8], which replaces α by α/P to account for the multiple tests; P is the number of individual z values being tested. If the standard deviation is not known but must be estimated, critical values from the student's t distribution are used.

The second multivariate fault detection technique uses a multivariate statistical distance, the Mahalanobis distance, to jointly monitor P measured variables simultaneously. The Mahalanobis distance of vector-valued residual $\mathbf{r}_i$ from the mean or target vector $\bar{\mathbf{r}}$ is

$$Md_i = (\mathbf{r}_i - \bar{\mathbf{r}})' \Sigma^{-1} (\mathbf{r}_i - \bar{\mathbf{r}}) \quad (3)$$

if the covariance is known. The Mahalanobis distance for $\mathbf{r}_i$ is compared to user-specified critical values from either the chi-squared distribution (or the weighted F distribution if Σ is not known but must be estimated). In the present application, the target vector $\bar{\mathbf{r}}$ is zero.

The covariance matrix Σ is necessary for calculation of the test statistics, $z_{ip}(\alpha)$ and Md_i. For the three-tank system, this was obtained by performing one thousand 10-h simulations under no-fault operating conditions. The flow rates and initial tank conditions were the same for each simulation except for the application of randomly distributed uncertainty to all measured variables. For the initial tank conditions, the uncertainties were applied once at the beginning of each simulation. For the flow rates, we assumed that new measured values became available every 0.1 h at which time the model values were updated. Model predictions at the end of each 10-h simulation were compared to measured values obtained at the same time and residuals calculated.

We also performed a principal component analysis of the 1000 simulated residuals. Principal components analysis is often used as a dimensionality reduction method when correlated variables are present. For data vectors containing P elements, it may be assumed that the components corresponding to the Q ($Q \leq P$) largest eigenvalues explain the important internal structure of the data. We used an approach presented in a paper on multivariate process control by Jackson [9]. Jackson suggested rescaling the eigenvectors by the eigenvalues so that each score has a mean of zero and a variance of one. Thus the score, t_{ij} for residual i and principal component j can be tested directly against critical values from the N(0,1) distribution to determine if the particular principal component score may be an outlier. In addition the Mahalanobis distances are easily calculated directly from the scaled scores. For observation i this is

$$Md_i = \mathbf{t}_i \, \mathbf{t}_i' . \qquad (4)$$

We investigated two diversion scenarios: 1) a steady leak from the second tank without replacement and 2) the same leak but the lost solution replaced with water. In practice, the model would not know about a leak and thus would make *erroneous* predictions because it assumes "normal" operations. The true conditions, i.e., the loss because of leakage, are reflected in the measured data. Results of fault detection tests for a 0.5-L/h leak, which was easily detected, are summarized in Fig. 4. Concentrations of plutonium and uranium plus density were detected as outliers in the second tank (Fig. 4b) under the leak-with-water-replacement scenario whereas only volume (Fig 4a) was detected as a fault for the leak-without-replacement scenario. Replacing the removed volume with water diluted the concentrations enough to make a large difference in all concentration variables as well as in density, which is based on concentrations. The Mahalanobis distances are shown in Fig. 5 for three different leak rates. In all cases the values are larger, thus more statistically significant, for the leak with replacement scenario. The 0.5-L/h rate was the only one significant at the 5% level.

3. UNIVARIATE TIME SERIES ANALYSIS

We considered univariate time series such as MUF values or other statistics arising from safeguards.

In computer simulations we experimented with several nonlinear estimation methods using both linear and nonlinear simulated data sets. Our approach assumes that the same functional dependence between an observation and some subset of the previous observations holds throughout the entire time series. If this assumption is not valid, the time series must be divided into subsets in which the assumption is satisfactory. This requires detailed knowledge of the process that is generating the sequence. Except for this potentially serious problem, the implementation of our procedures is straightforward.

We implemented FORTRAN computer codes to perform the two main activities in estimating the conditional mean: choosing the degree of smoothing and estimating the lag, i.e., number of previous observations directly affecting the present observation.

Empirical estimates of the lag may enhance understanding of the processes generating the data. In many cases, we have a good idea of a value for the lag. For example, in ordinary MUF sequences, a value of one is often a good first approximation for the lag. This is because the ending inventory for MUF_{j-1} is the beginning inventory for MUF_j. Therefore, if we ignore the effects of systematic errors, the lag is one.

Our main goal is better detection of anomalies through use of the best techniques for predicting future values of the time series. By best, we mean that the standard error of residuals (MSEP) is minimized.

To illustrate nonlinear modeling, we compared the MSEP using two linear and three nonlinear estimation methods as follows:

(1) Divide the time series vector into testing and training sets and assume no loss has occurred.

(2) Use the training set to estimate the conditional mean making either no assumption about the functional form of the conditional mean or assuming that the conditional mean is linear.

(3) Compare the MSEP evaluated in the test set for the linear and the nonlinear methods.

Regarding the MSEP, we have analyzed simulated data from the following six time sequences, each observed with error.

(a) $x_t = 1 - 1.4x_{t-1}^2 + 0.3x_{t-2}$ (nonlinear),

(b) $x_t = 4x_{t-1}(1 - x_{t-1})$ (nonlinear),

(c) $X_t = a_0 + b_1e_{t-1} + b_2e_{t-2} + e_t$ (linear),

(d) $X_t = a_0 + a_1X_{t-1} + a_2X_{t-2} + e_t$ (linear),

(e) $X_t = a_0 + b_1e_{t-1} + e_t$ (linear), and

(f) $X_t = a_0 + a_1X_{t-1} + e_t$ (linear).

We have used lower case for the two nonlinear time series, (a) and (b), because we generated the data deterministically and then added observational error. Specifically, we generated the x_t's and then added independent $N(0,0.05^2)$ random variables to represent observational errors. For the linear series, the errors were independent N(0,1) random variables. Therefore, the theoretically lowest achievable MSEP is 0.0025 for series (a) and (b), and 1 for series (c-f). For each of the six time series, we generated training vectors with 1000 observations and testing vectors with 1000 observations. The MSEPs for two methods of linear estimation and for three types of nonlinear estimation are shown in Table I for the six cases. We include the sample variance s^2 for each case because s^2 would be the MSEP if we used the sample mean as the predicted value.

In Table I, the first linear method is denoted Linear1, the second linear method is denoted Linear2, and similarly for the three nonlinear methods. The first linear method fits the best possible autoregressive moving average (ARMA) model to the observed sequence. The second linear method fits a linear model to the regression of X_t on X_{t-1} or on X_{t-1} and X_{t-2}, depending on which gives a better fit.

Table I. Simulation Results for Mean Square Error of Prediction (MSEP)

Time Series	s^2	Linear1	Linear2	Nonlinear1	Nonlinear2	Nonlinear3
a	0.53	0.47	0.46	0.06	0.01	1.21
b	0.13	0.13	0.13	0.02	0.10	0.94
c	1.36	0.98	1.14	1.48	0.821	0.37
d	1.68	0.98	0.99	1.05	0.60	0.84
e	1.14	1.01	0.99	1.01	0.87	1.48
f	1.30	1.00	0.98	1.00	0.76	0.80

The first nonlinear method is a conditional mean estimator. For the lag = 1 case, our estimator is

$$\hat{M}(x) = \left[\frac{1}{n-1}\sum_{j=1}^{n-1} X_{j+1}\kappa\left(\frac{x-X_j}{h}\right)\right] \Bigg/ \left[\frac{1}{n}\sum_{j=1}^{n}\kappa\left(\frac{x-X_j}{h}\right)\right] , \tag{5}$$

where κ is called the kernel. It is usually assumed that $\kappa(x)$ evaluated at $x = 0$ is the maximum value of κ, and that $\kappa(x)$ is a decreasing function of $|x|$. It is usually further assumed that κ is a symmetric probability density function such as the standard normal density. The parameter h is the bandwidth, which determines the amount of smoothing. For more detail, see Ref. [5] but the idea in Eq. (5) is straightforward. We have n observations, $X_1, X_2, \cdots, X_n$, and seek an estimate of X_{n+1} given the value $X_t = x$. The idea is to use all of the first $n-1$ observations, but weigh most heavily the observations that are most near the value x. The extension to higher lags is straightforward. We present results here only for the lag one case for the first nonlinear method. The second method is different from the first method in that the second method does attempt to estimate the lag and uses a different method to choose the bandwidth. Using the second method, the best estimate of the lags for time series (a) through (f) was $d = 5, 2, 2, 2, 2,$ and 1, respectively, whereas the correct lags are 2, 1, 2, 2, 1, and 1. The third nonlinear method is a computationally intensive method which appears to be an inconsistent performer in our experiments to date. The method uses the k nearest neighbors of each point to fit a local linear model at that point. The overall model is then piecewise linear, but can be made to look rather smooth if the pieces are sufficiently short.

In Table I there is not much difference in MSEP between the two linear methods except for series c where Linear1 gives a lower value. We expect the second linear method to perform worse on MA models than the first linear method because the second linear method relies on the true time series being an autoregressive series. Similarly, all three of our nonlinear methods are designed for autoregressive series. However, it is possible to extend our nonlinear methods to accommodate MA models. The details of this extension can be found in Ref. [10].

In comparing the linear methods to the nonlinear methods, note that cases a and b are the only nonlinear series. For series a and b there is a clear advantage in using the Nonlinear1 over the linear methods. However, notice that the other two nonlinear methods do not perform consistently. In fact, because the error variance for the two linear time series was 1.0, the theoretically lowest achievable variance for predicting them is 1.0. Therefore, Nonlinear2 and 3 sometimes give misleadingly low estimates of the true MSEP. We currently have no explanation for this behavior. At present we prefer the relative simplicity of Nonlinear1 and are pleased with its performance on both linear and nonlinear time series.

4. SUMMARY AND CONCLUSIONS

Both the multivariate process fault detection and the nonlinear time series methods are fairly easy to implement. With respect to safeguards, the main issue is whether international inspectors will be granted access to the larger amounts of data expected from modern reprocessing plants. Until such data are available, we can only test our methods on simulated data.

For the three-tank problem, univariate tests on individual variables as well as on individual principal components were equally effective at detecting losses of material. Because the principal components are linear combinations of individual variables, they might be expected to provide more sensitive detection of outliers or faults for situations where two or more correlated variables are affected by a fault. With this simulation, the correlations were not strong enough to observe this effect. Multivariate tests based on the Mahalanobis distances were never as sensitive as the univariate tests probably because the sensitivity is diminished somewhat by those variables not affected by a leak. The leak with replacement scenario was detected with slightly more sensitivity than was the leak without replacement perhaps because replacing the lost volume with water affects four variables (density and the concentrations of nitric acid, plutonium, and uranium) rather than just one as does a leak without replacement.

For univariate time series, our current recommendation is to include techniques that can test for nonlinearity in a package of evaluation methods for time series. If tests do not indicate nonlinearity, there is no need to apply nonlinear estimation methods. If tests do indicate nonlinearity, we recommend using nonlinear techniques for estimating the expected value of an observation in the time series sequence. Presumably, if a time series sequence fails the tests for linearity, future expected values will be a nonlinear function of some subset of the previous observations. Another potential advantage of nonlinear modeling could result from an improved understanding of mechanisms generating the data through detection of unexpected functional dependencies. Thus, for reasons other than anomaly detection, we may wish to analyze time sequences containing many elements using nonlinear methods should they become available in the future.

Judging from results obtained using simulated data, the multivariate process control methods will improve our ability to detect loss by exploiting the redundancies between modeled and measured variables. The nonlinear time series methods perform as well as linear methods when the true functional form is linear and outperform the linear methods when the true functional form is nonlinear. However, the price paid in both cases is that more measurements must be made.

REFERENCES

[1] FRANK, PAUL M., Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy—A survey and some new results, Automatica **26** (1990) 459-474.

[2] BURR, T.L. and WANGEN, L.E., Multivariate Diagnostics and Anomaly Detection for Nuclear Safeguards, Los Alamos National Laboratory (1994) (report in preparation).

[3] CHATFIELD, C, The Analysis of Time Series, Chapman and Hill, New York (1989).

[4] TONG, H., Nonlinear Time Series, Clarendon Press, Oxford (1990).

[5] AUESTAD, BJORN and TJOSTHEIM, DAG, Identification of nonlinear time series: First order characterization and order determination, Biometrika **77** 4, (1990) 669-687.

[6] BEYERLEIN, A., Clemson University, private communication, 1992.

[7] WATSON, W.B. and RAINEY, R.H., Modification of the SEPHIS Computer Code for Calculating the Purex Solvent Extraction System, ORNL-TM-5123, Oak Ridge National Laboratory (December 1975).

[8] JOHNSON, R.A. and WICHERN, D.W., Applied Multivariate Statistical Analysis, Prentice-Hall, Englewood Cliffs, New Jersey (1988)

[9] JACKSON, J. EDWARD, Multivariate quality control communications in statistics, Theoretical Methods **14** (1985) 2657-2688.

[10] BURR, T.L., Time Series Primer, N-4/93-1263, Los Alamos National Laboratory (December 1993).

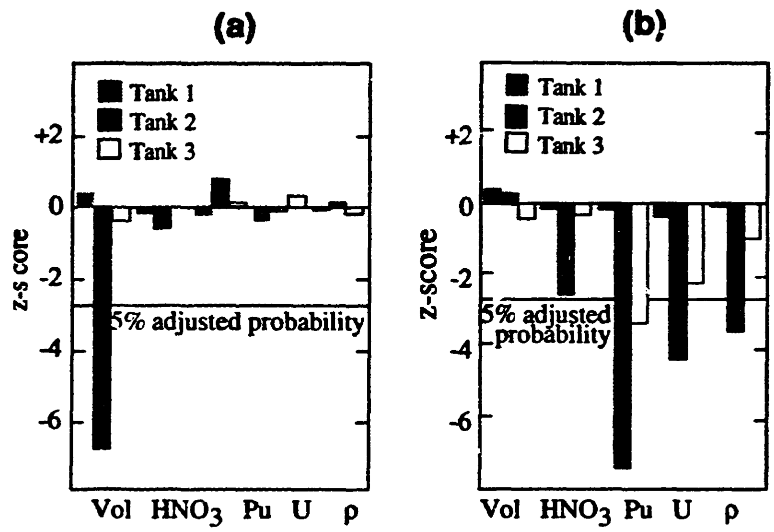

Fig. 4. Z scores for each variable in the three-tank system for a leak rate of 0.5 L/h. (a) is without replacement of lost volume. In b the lost volume is replaced with water.

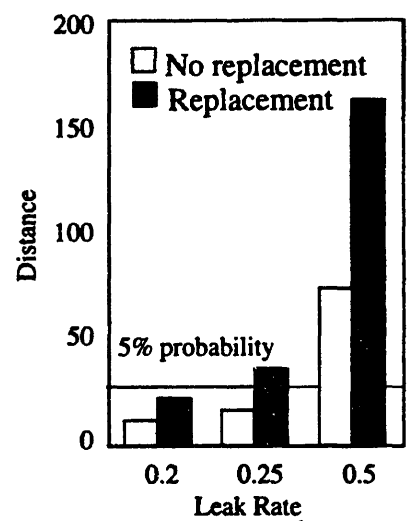

Fig. 5. Mahalanobis distances for three leak rates under replacement and no replacement scenarios.

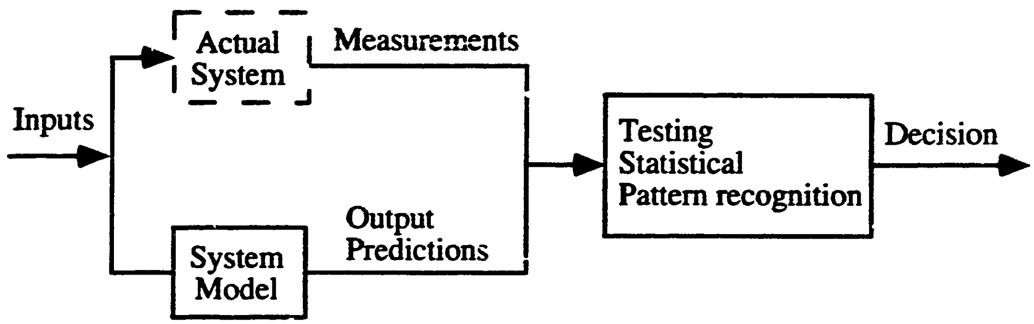

Fig. 1. Process fault detection.

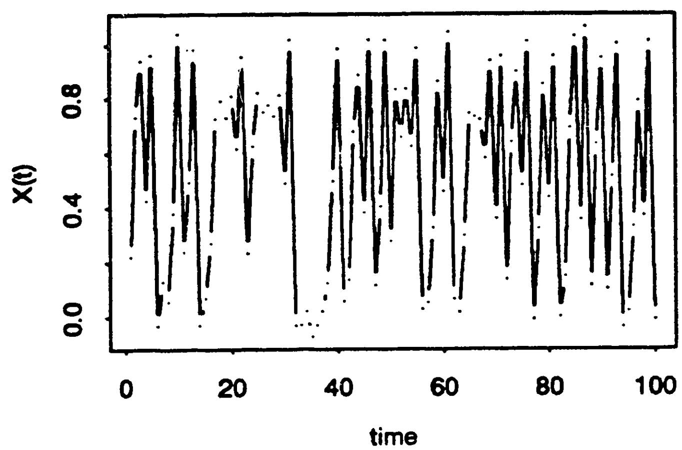

Fig. 2a. Time-series plot of 100 observations of a nonlinear time series.

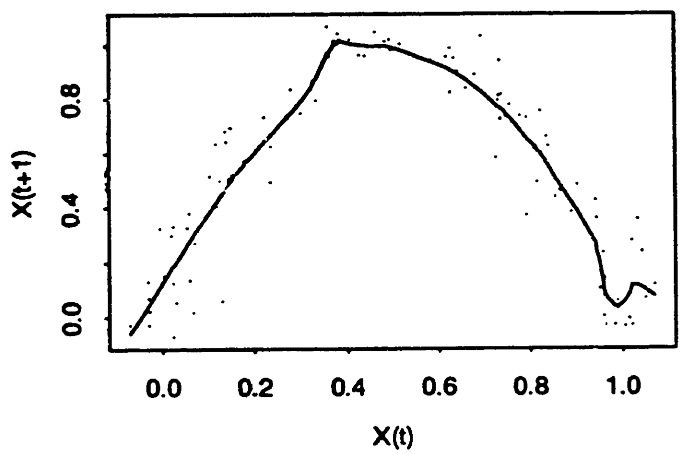

Fig. 2b. Same time-series as in (a) with nonlinear fit superimposed on scatterplot of $X(t + 1)$ versus $X(t)$.

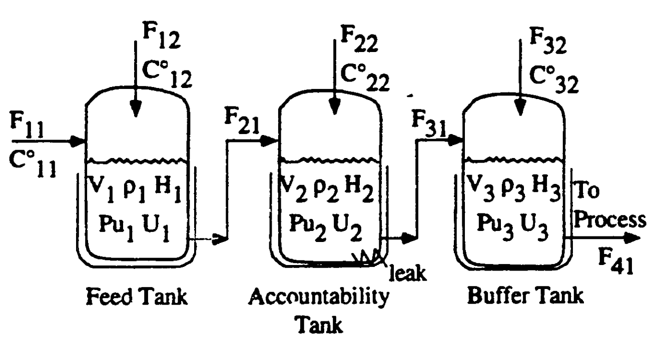

Fig. 3. A three-tank system. V, P. H, Pu, and U refer to volumes, densities, and concentrations of nitric acid, plutonium, and uranium. C°s refer to concentrations in input flows Fs.

DATE
FILMED
3/8/94
END