RECENT ADVANCES IN MULTIVARIATE STATISTICAL PROCESS CONTROL: IMPROVING ROBUSTNESS AND SENSITIVITY

B. M. Wise and N. L. Ricker

Center for Process Analytical Chemistry and Department of Chemical Engineering BF-10, University of Washington, Seattle, Washington 98195

Abstract. Several extensions are made to the theory of multivariate process monitoring *via* Principal Components Analysis (PCA). An important robustness issue is addressed: the continued use of the PCA model after detection of a sensor failure. Without some adjustment, a single failed sensor can obscure other failures, thus rendering the monitoring method useless. It is shown here that one can calculate an estimate of the output of the failed sensor that is most consistent with the PCA model of the process. This estimate allows continued use of the model. Under some circumstances, replacing the failed output with this estimate is equivalent to rebuilding the entire PCA model. Partial Least Squares (PLS) regression can be used in a manner similar to PCA for process monitoring. It is shown that PLS is fundamentally more sensitive to sensor failures than PCA. Unlike PCA, however, the PLS monitoring scheme maps state information into the model residuals. For this reason, changes in the process state covariance and autocovariance can invalidate calculated PLS model residual limits. The failed sensor problem is also solved for the PLS monitoring method.

Keywords. Principal components; partial least squares; failure detection; multivariable systems.

INTRODUCTION

Principal Components Analysis (PCA) has shown great utility as a Multivariate Statistical Process Control (MSPC) tool when the samples can be considered independent observations (Jackson, 1981; Veltkamp and co-workers, 1990; Wise and co-workers, 1991). The utility of the method for monitoring dynamic processes has also been shown (Kresta, MacGregor and Marlin, 1990; MacGregor, 1989; Wise and Ricker, 1989). Recently, a theoretical basis for the use of PCA in dynamic systems with measurement redundancy has been developed (Wise and co-workers, 1990). Briefly, the application of PCA to the measurements from such systems confines variations in the process states to the PCA "scores", and under normal conditions the PCA residuals are white. Statistical tests can be applied to the PCA residuals to detect and diagnose a variety of abnormal events. A key advantage of this approach is that a complete dynamic model of the system need not be developed.

Some important issues that bear on the robustness and sensitivity of the monitoring method have yet to be addressed, however. One of particular interest concerns the continued use of the PCA model after a sensor has been identified as having failed. If the data from the failed sensor continues to be included in the monitoring system it has the tendency to “mask” subsequent abnormal events.

It is also known that the PCA monitoring method may be quite insensitive to particular sensor failures. It has been proposed (Wise, Ricker, and Veltkamp, 1989) that collections of Partial Least Squares (PLS) regression models can be used in a manner similar to PCA for process monitoring. Results have shown that this method may be more sensitive than PCA, but the reason for this increased sensitivity was not clarified in previous work, and the potential pitfalls of implementing this PLS approach have not been addressed.

It is the purpose of this work to address both the robustness and sensitivity issues. In the sections that follow a method for estimating the outputs of failed sensors is proposed and it is shown that, under certain conditions, this method is equivalent to rebuilding the entire PCA model. This allows continued use of the original model. The relationship between PCA and PLS monitoring is also explored and the reasons for the increased sensitivity of the PLS method are demonstrated. Potential problems with PLS monitoring are also investigated. Finally, the failed sensor estimation problem is solved for PLS models.

BACKGROUND

The PCA method and the basis for its use in process monitoring is reviewed in the following sections. We also review the proposed method for generation of PCA-like models *via* PLS.

Principal Components Analysis

Let $\mathbf{X}$ be an m by n data matrix in which the rows are samples and the columns are variables. We assume that $\mathbf{X}$ has mean-

centered such that its columns all have zero mean. In PCA, $\mathbf{X}$ is decomposed into the sum of the product of *n* pairs of vectors (Jackson, 1981; Veltkamp and co-workers, 1990). Each pair consists of a *n* by *1* vector called the *loadings*, $\mathbf{p}_i$, and a *m* by *1* vector called the *scores*, $\mathbf{t}_i$. Thus $\mathbf{X}$ can be written as

$$\mathbf{X} = \mathbf{t}_1\mathbf{p}_1^T + \mathbf{t}_2\mathbf{p}_2^T + \ldots + \mathbf{t}_n\mathbf{p}_n^T \quad (1)$$

The matrix of loadings vectors $\mathbf{P}$ forms a new orthogonal basis for the space spanned by $\mathbf{X}$ and the individual $\mathbf{p}_i$ are the eigenvectors of the covariance matrix of the mean-centered data matrix $\mathbf{X}$, defined as:

$$\text{covariance}(\mathbf{X}) = \frac{1}{m - 1}\left(\mathbf{X}^T\mathbf{X}\right) \quad (2)$$

Thus

$$\text{covariance}(\mathbf{X})\mathbf{p}_i = \lambda_i\mathbf{p}_i \quad (3)$$

where λ_i is the eigenvalue associated with the eigenvector $\mathbf{p}_i$. If $\mathbf{X}$ has been *autoscaled* (*i.e.*, mean-centered with each column scaled to unit variance) the covariance matrix becomes the correlation matrix. The loadings vectors $\mathbf{p}_i$ are referred to as principal components because they are linear combinations of the original variables that together account for large fractions of the variance in the original matrix. Each of the scores vectors $\mathbf{t}_i$ is simply the projection of $\mathbf{X}$ onto the new basis vector $\mathbf{p}_i$:

$$\mathbf{t}_i = \mathbf{X}\mathbf{p}_i \quad (4)$$

PCA is very closely related to the Singular Value Decomposition (SVD) (Strang, 1980) where $\mathbf{X}$ may be decomposed as

$$\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T \quad (5)$$

where $\mathbf{V}$ contains the eigenvectors ($\mathbf{p}_i$) and $\mathbf{S}$ is a diagonal matrix containing the square roots of the eigenvalues (the singular values) of the covariance matrix of $\mathbf{X}$.

When PCA is done on a data set, it is often found (and it is generally the objective) that only the first few eigenvectors are associated with systematic variation in the data, and the remaining eigenvectors are associated with “noise”. Noise in this case refers to uncontrolled experimental and instrumental variations arising from random processes. PCA models are formed by retaining only the eigenvectors that are descriptive of systematic variation in the data. Determination of the proper number of eigenvectors can be done by cross-validation or other techniques (Malinowski, 1977, 1987). Once the PCA model is formed, new data can be viewed as projections onto single eigenvectors (scores plots) or the plane formed by pairs of eigenvectors. The scores can be used to obtain the "PCA filtered estimate" of a given sample, *i.e.* the projection of the sample into the PCA model. For a reduced order model, $\mathbf{P}_k$, (where only

the first k of the n total eigenvectors are retained) and a new sample, $\mathbf{x}_i$, this is obtained from:

$$\hat{\mathbf{x}}_i = \mathbf{t}_{ki}\mathbf{P}_k^T = \mathbf{x}_i\mathbf{P}_k\mathbf{P}_k^T \quad (6)$$

where $\mathbf{t}_{ki}$ is the (1 by k) vector of scores on the model $\mathbf{P}_k$ for sample $\mathbf{x}_i$.

The "goodness" of fit between new data and the model can be monitored by calculating residuals, *i.e.*, the difference between a sample and its PCA estimate. The residual $\mathbf{r}_i$ for sample $\mathbf{x}_i$ is given by

$$\mathbf{r}_i = \mathbf{x}_i - \hat{\mathbf{x}}_i = \mathbf{x}_i(\mathbf{I} - \mathbf{P}_k\mathbf{P}_k^T) \quad (7)$$

The magnitude of the residual for any sample $\mathbf{x}_i$ is

$$Q = \mathbf{r}_i\mathbf{r}_i^T = \mathbf{x}_i(\mathbf{I} - \mathbf{P}_k\mathbf{P}_k^T)\mathbf{x}_i^T \quad (8)$$

Basis for PCA Monitoring

The discrete state-space process model has the form:

$$\mathbf{x}(k+1) = \Phi\mathbf{x}(k) + \Gamma\mathbf{u}(k) + \nu(k) \quad (9)$$

$$\mathbf{y}(k) = \mathbf{C}\mathbf{x}(k) + \mathbf{D}\mathbf{u}(k) + \mathbf{e}(k) \quad (10)$$

Assuming that the process has n states, r inputs and p measurements, then $\mathbf{x}(k)$ is the (n by 1) vector of state variables at time k, $\mathbf{u}(k)$ is the (r by 1) vector of process inputs, $\nu(k)$ is the (n by 1) vector of state disturbances, $\mathbf{y}(k)$ is the (p by 1) vector of process measurements $\mathbf{e}(k)$ is the vector of measurement noise, and Φ, Γ, $\mathbf{C}$, and $\mathbf{D}$ are matrices of appropriate size. For most processes $\mathbf{D}$ is zero; process inputs rarely have an instantaneous effect on the process outputs. The dimension of the state space can be greater than, less than or equal to the dimension of the measurement space. Many references concerning the state-space formalism are available (Kwakernaak and Sivan, 1972; Sage and White, 1977).

We have shown previously (Wise and co-workers, 1990) that, for processes with more measurements than states, proper application of PCA can facilitate the process monitoring and fault detection problem. This is done by identification of a PCA model of the outputs $\mathbf{y}(k)$. Under most circumstances, the PCA model will span the same space as the $\mathbf{C}$ matrix in the corresponding state-space representation of the process. When the PCA model is an accurate estimate of $\mathbf{C}$, variations in the process states $\mathbf{x}(k)$ appear primarily as variations in the PCA scores, while noise $\mathbf{e}(k)$ mainly affects the residuals. This allows one to consider only the noise properties when deriving statistical limits for the PCA residuals. In particular, the process dynamics need not be considered explicitly. This implies that conventional statistical methods that rely on independence of the

samples (such as t- and F-tests and the multivariate T^2) can be applied to the residuals.

PLS Monitoring

Many references are available on the PLS method (Geladi and Kowalski, 1986; Lorber, Wangen and Kowalski, 1987) so the algorithm will not be repeated here. Briefly, PLS works by simultaneously decomposing both the input data block **X** and the output data block **Y** (a vector if there is only one output variable) in such a way that the factor scores in the **X** and **Y** blocks have the maximum covariance. The number of factors or latent variables retained in the PLS regression model is optimized based on prediction through a series of cross-validations. The parameters used in PLS prediction can also be reduced to a single linear equation:

$$\hat{\mathbf{Y}} = \mathbf{XB} \tag{11}$$

where **B** is a matrix in the general case and a vector in the case of only one variable in the **Y** block. The result is Eq. (12) where k is the number of latent variables to be used in the prediction and it is assumed that the value of the term in brackets is equal to **I** for the case of i = 1.

$$\hat{\mathbf{Y}} = \mathbf{X} \sum_{i=1}^{k} b_i \left[\prod_{j=1}^{i-1} \left(\mathbf{I} - \mathbf{w}_j \mathbf{p}_j^T \right) \right] \mathbf{w}_i \mathbf{q}_i^T \tag{12}$$

In (12) the b_i are the inner relation coefficients, the $\mathbf{w}_j$ and the $\mathbf{p}_j$ are the **X** block weights and loadings and the $\mathbf{q}_i$ are the **Y** block loadings. Here the notation of Geladi has been used.

It is proposed that PLS can also be used to monitor processes in a fashion similar to the use of PCA models. This requires that PLS models be obtained that relate each process output to the remaining outputs in the system. Thus for a system with *n* outputs, *n* PLS models would be required. Fortunately, using the relationship given in Eq. (12) the *n* PLS models can be formed into a single matrix, with each model being a column vector. Because each of the variables does not contribute to its own prediction, the resulting prediction matrix, $\mathbf{M}_p$, has zeros on the diagonal. Thus the "PLS filtered estimate" of a data matrix **X** can be obtained by simple matrix multiplication

$$\hat{\mathbf{X}} = \mathbf{X}\,\mathbf{M}_p \tag{13}$$

A residuals matrix, $\mathbf{D}_{pls}$, can be calculated from

$$\mathbf{D}_{pls} = \mathbf{X} - \hat{\mathbf{X}} = \mathbf{X} - \mathbf{X}\,\mathbf{M}_p = \mathbf{X}(\mathbf{I} - \mathbf{M}_p) = \mathbf{X}\,\mathbf{R}_{pls} \tag{14}$$

where $\mathbf{R}_{pls}$ is the PLS equivalent of the $\mathbf{I} - \mathbf{PP}^T$ matrix in PCA. Once the $\mathbf{R}_{pls}$ matrix has been calculated it can be used with new data to produce PCA-like residuals. These residuals can be used with statistical tests to detect any changes to the process or its

sensors. As will be shown, however, there are some important differences between PCA and PLS monitoring.

IMPROVING ROBUSTNESS

As mentioned previously, if the data from failed sensors continues to be included in the PCA monitoring system it can mask additional failures, rendering the monitoring system useless. In this sense the system is not very robust; it is only capable of detecting the first sensor failure. It seems logical that there should be an optimal way of either replacing "bad" data or modifying the existing PCA model so that changes in other variables can be observed, without reforming the entire model. The solution to the problem of replacing data from single or multiple sensors follows.

Estimating Failed Sensor Outputs

Assume for the moment that the PCA model of the process of interest has been calculated, and that the matrix used for calculating the residuals has been obtained:

$$\mathbf{I} - \mathbf{P}_k\mathbf{P}_k^T = \mathbf{R}_{pca} \tag{15}$$

The Q residual for any sample $\mathbf{x}$ (a row vector) can then be calculated as:

$$Q = \mathbf{x}\mathbf{R}_{pca}\mathbf{x}^T \tag{16}$$

Suppose now that one or more sensors have failed and have been detected. Further, suppose that it is convenient to partition $\mathbf{x}$ (possibly by rearranging the columns of $\mathbf{x}$) into a group of "bad" sensors $\mathbf{x}_b$, and a group of good sensors $\mathbf{x}_g$. Thus:

$$\mathbf{x} = [\mathbf{x}_b\ \mathbf{x}_g] \tag{17}$$

It is possible to partition $\mathbf{R}$ into parts that act on each of the groups of good and bad sensors individually:

$$\mathbf{R}_{pca} = \begin{bmatrix} \mathbf{R}_{11} & \mathbf{R}_{12} \\ \mathbf{R}_{21} & \mathbf{R}_{22} \end{bmatrix} \tag{18}$$

The Q residual is now calculated as:

$$Q = \begin{bmatrix} \mathbf{x}_b\ \mathbf{x}_g \end{bmatrix} \begin{bmatrix} \mathbf{R}_{11} & \mathbf{R}_{12} \\ \mathbf{R}_{21} & \mathbf{R}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x}_b^T \\ \mathbf{x}_g^T \end{bmatrix} \tag{19}$$

By multiplying this through, and using the fact that $\mathbf{R}_{21} = \mathbf{R}_{12}^T$, the following expression for Q is obtained:

$$Q = \mathbf{x}_b\mathbf{R}_{11}\mathbf{x}_b^T + \mathbf{x}_g\mathbf{R}_{21}\mathbf{x}_b^T + \mathbf{x}_b\mathbf{R}_{21}^T\mathbf{x}_g^T + \mathbf{x}_g\mathbf{R}_{22}\mathbf{x}_g^T \tag{20}$$

It would seem logical to estimate $\mathbf{x}_b$ based on the value which minimizes Q. This estimates the failed sensor output as the value which is most consistent with the PCA model. Note that in Eq. (20) the last term is a function of $\mathbf{x}_g$ and $\mathbf{R}_{22}$ only, and is

therefore fixed. The problem of minimizing Q then becomes the problem of minimizing the first three terms on the right hand side of (20). Thus the objective becomes to find:

$$\hat{\mathbf{x}}_b \quad \{\mathbf{x}_b\mathbf{R}_{11}\mathbf{x}_b^T + \mathbf{x}_g\mathbf{R}_{21}\mathbf{x}_b^T + \mathbf{x}_b\mathbf{R}_{21}^T\mathbf{x}_g^T\} = \min \quad (21)$$

Fortunately, this problem can be easily solved by "completing the squares" (Åstrom and Wittenmark, 1984). The result is:

$$\hat{\mathbf{x}}_b = -\mathbf{x}_g\mathbf{R}_{21}\mathbf{R}_{11}^{-1} \quad (22)$$

Because $\mathbf{R}_m$ is fixed, the value of $\mathbf{x}_b$ for any new sample can be calculated from

$$\hat{\mathbf{x}}_b = \mathbf{x}_g\mathbf{R}_R \quad (23)$$

where $\mathbf{R}_R$ is the regression matrix formed from $\mathbf{R}_m$:

$$\mathbf{R}_R = -\mathbf{R}_{21}\mathbf{R}_{11}^{-1} \quad (24)$$

which can be calculated once and retained. Calculation of $\mathbf{R}_R$ should be possible because $\mathbf{R}_{11}$ will be positive definite provided that none of the variables to be solved for are totally independent, *i.e.*, not correlated with any of the other outputs of the system. Computationally, the biggest problem occurs when the bad variables are not in a "convenient" location and the matrices $\mathbf{R}_{11}$ and $\mathbf{R}_{21}$ must be extracted from $\mathbf{R}_{pca}$. The process for doing this is shown in Fig. 1. Here sensors 2 and 5 are assumed to have failed. The figure shows how the parts of the original $\mathbf{R}_m$ matrix map into $\mathbf{R}_{11}$ and $\mathbf{R}_{21}$.

Once $\mathbf{R}_R$ has been obtained, it can be mapped into an identity matrix in such a way that multiplication of new samples by this matrix, the "replacement" matrix $\mathbf{R}_m$, results in the bad variables being replaced by their values that minimize Q. The mapping is shown in Fig. 2. The solid black squares in the figure represent "ones", the white portions represent "zeros" and shaded portions are occupied by the regression matrix.

It is found that when new samples are multiplied by $\mathbf{R}_m$ the PCA residuals of the replaced variables are identically zero. This can be seen by substituting the estimates of $\mathbf{x}_b$ back into Eq. (20). This also leads to an expression for the minimum value of Q.

$$Q_{min} = -\mathbf{x}_g\mathbf{R}_{21}\mathbf{R}_{11}^{-1}\mathbf{R}_{21}^T\mathbf{x}_g^T + \mathbf{x}_g\mathbf{R}_{22}\mathbf{x}_g^T \quad (25)$$

The problem now is that the resulting Q values will tend to be artificially low. In order to remedy this, it would be possible to add a "white" noise signal of zero mean and appropriate variance to the estimate of the failed sensor outputs. However, it may be that for most practical monitoring problems adding noise is unnecessary. Using the original calculated limits may be close enough, particularly in cases where there are many variables

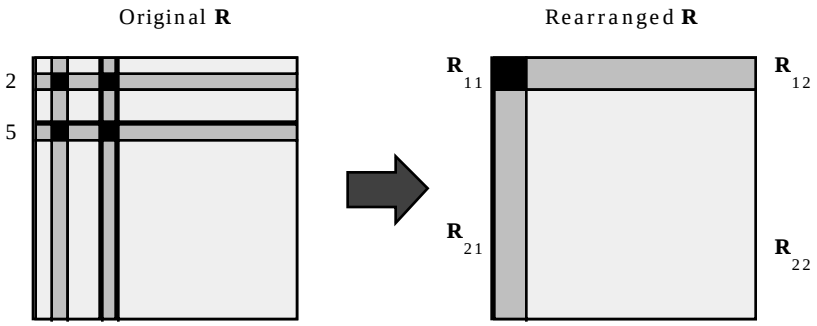

Fig. 1. Obtaining $\mathbf{R}_{11}$ and $\mathbf{R}_{21}$ from original $\mathbf{R}_{pca}$ matrix.

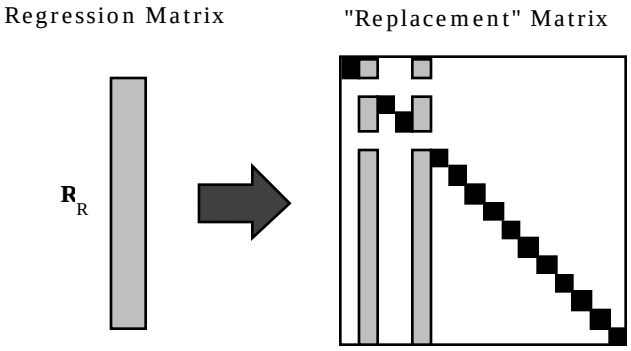

Fig. 2. Mapping of $\mathbf{R}_R$ into $\mathbf{R}_m$.

relative to the number of PCs retained in the model. In this case the loss of the variance due to one variable has only a small affect on the total residual. The variance of other residuals will be slightly decreased, making the original limits at an effectively higher confidence level. In many applications the difference would be insignificant.

The alternative to replacing bad variables in an existing PCA model is to re-calibrate the model, omitting the failed sensor. This approach was compared to the method outlined above. The somewhat surprising result is that the two methods are equivalent. In the noise free case, where the data is truly rank deficient, the two methods produce identical results (Wise, 1991). In the presence of noise, the solutions are approximately equal and usually very close to one another, provided that a sufficient number of samples are available for re-calibration.

Example of Failed Sensor Estimation

The system given in Wise and co-workers (1990) is used here to generate an example of improved robustness. Briefly, the system model has 5 states, 5 inputs and 10 measurements. The system is driven by white noise of unit variance and the variance of the measurement noise is equal to the deterministic variance of the outputs. A 5 PC model was identified from 1000 samples. A new data set of 500 samples was generated, but this time a bias of 3 units was added to output 1 starting at sample #201 and to output 5 at sample #301. The upper plot in Fig. 3 shows Q for the new data set along with the 95% Q limit. Note how the system detects a change at sample #201, as expected. The change at sample 301, however, is not obvious. In the bottom

plot the bad output was replaced as above starting at sample #201. Note how the failure at sample #301 is now obvious.

In Fig. 4 the residuals resulting from replacing the failed output are compared to the residuals on an entirely new model constructed without the failed sensor. Residuals are shown in the upper plot for sample #201 and in the lower plot for sample #301. Note how the replaced output residuals (circles) are zero on the first output, as expected. The new model residuals (x's) are nearly identical to the replaced output residuals for outputs 2 to 10. (There is no output 1 in the new model.)

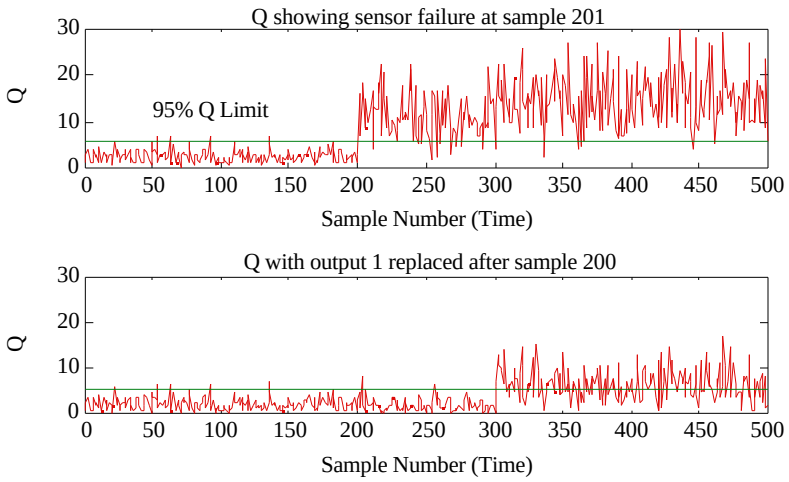

Fig. 3. Q with and without replacing first failed sensor output.

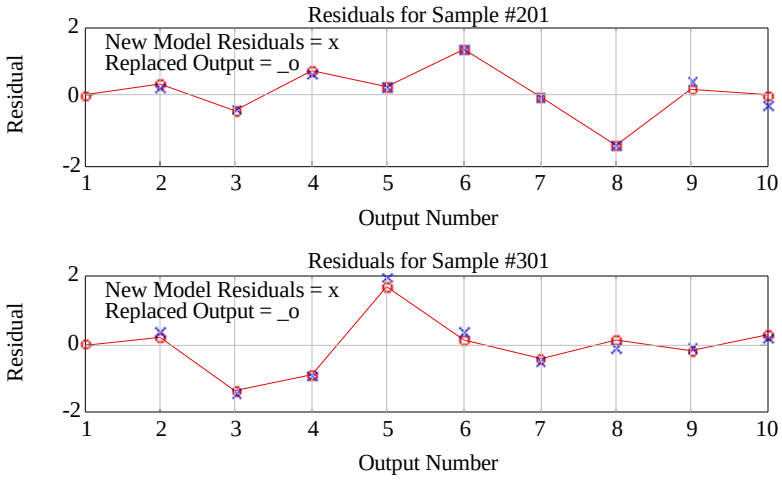

Fig. 4. Comparison of residuals with replaced output #1 and residuals on new model for sample #201 and #301.

This example demonstrates that estimating the outputs from failed sensor is a viable alternative to rebuilding the entire PCA model. One advantage of the method is that it allows continued use of any plots being used for monitoring, including any scores plots, without modification. Detection of additional failures is also enabled, though the original statistics on the residuals will be only approximate.

PLS MONITORING

In Wise, Ricker and Veltkamp (1989) it was demonstrated that PCA-like residuals could be generated with PLS, and that the detection limits were often much improved over PCA monitoring. This section provides a basis for the PLS monitoring approach and compares the effectiveness of the PLS monitoring method to PCA monitoring using synthetic data. The section closes with the solution to the "bad variable" problem for PLS monitoring that was solved for analogous PCA problem in the previous section.

A Basis for PLS Monitoring

In the following paragraphs it will be shown that PLS monitoring is a logical extension of the PCA monitoring. The argument proceeds as follows: it is shown that a PCA model can be converted to another form that has detection power identical to the original model. This transformed model has the form of a collection of regression models, arranged in a manner similar to the collection of PLS models described above. A collection of PLS models, therefore, is simply an optimization of these regression models to improve their predictive ability. Furthermore, the PLS models make no assumptions concerning the intrinsic dimensionality of the output variable space.

Suppose that we have a discrete LTI process with noise and with fewer states than measurements in the state-space form as defined by Eqs. (9) and (10). Furthermore, suppose that an accurate PCA model $\mathbf{P}_k$ has been determined for the system, *i.e.*, $\mathbf{P}_k$ spans the same space as $\mathbf{C}$ from the "true" state-space model of the process. Using the notation of the previous section, with $\mathbf{R}_{pca}$ partitioned as in Eq. (18), it is possible to write the residual on the first variable $r_1(k)$ as

$$r_1(k) = \mathbf{R}_{11}\mathbf{C}_1\mathbf{x}(k) + \mathbf{R}_{12}\mathbf{C}_2\mathbf{x}(k) + \mathbf{R}_{11}e_1(k) + \mathbf{R}_{12}\mathbf{e}_2(k) \quad (26)$$

where $\mathbf{C}_1$ corresponds to the first row of the state-space $\mathbf{C}$ matrix and $\mathbf{C}_2$ is equal to the remaining rows 2 to p. It is known that if the PCA model is accurate, the first two terms on the right hand side of Eq. (26) sum to zero, *i.e.* the states make no contribution to the residuals. Thus

$$r_1(k) = \mathbf{R}_{11}e_1(k) + \mathbf{R}_{12}\mathbf{e}_2(k) \quad (27)$$

Suppose now that the PCA model is transformed so that the residuals on the transformed model $\mathbf{r}_t$ are defined as

$$\mathbf{r}_t(k) = \mathbf{y}(k) - \hat{\mathbf{y}}(k) \quad (28)$$

where each of the y_i is estimated based on the variable replacement method defined above, *e.g.*, for the first variable in the system

$$\hat{y}_1(k) = -\mathbf{y}_2\mathbf{R}_{12}\mathbf{R}_{11}^{-1} \quad (29)$$

The first transformed model residual $r_{t1}(k)$ can now be written in terms of the state-space model parameters as

$$r_{t1}(k) = \mathbf{C}_1\mathbf{x}(k)+e_1(k)+\mathbf{C}_2\mathbf{x}(k)\mathbf{R}_{21}\mathbf{R}_{11}^{-1}+\mathbf{e}_2(k)\mathbf{R}_{21}\mathbf{R}_{11}^{-1} \quad (30)$$

Note the correspondence between Eq. (30) and Eq. (26). It is easily seen that the residual $r_1(k)$ is different from $r_{t1}(k)$ by a factor of $\mathbf{R}_{11}^{-1}$. Furthermore, it can be seen that the transformed model will have fault detection "power" identical to the original model: the ratio of the expected size of the residuals to the size when a fault has occurred is the same for both the original and the transformed models. When the model is transformed in this manner, if a single variable changes, its residual changes by the same amount. Thus, the entire PCA model for calculating residuals $\mathbf{R}_{pca}$ can be transformed to $\mathbf{R}_t$ as follows:

$$\mathbf{R}_t = \mathbf{R}_{pca}(\text{diag}(\mathbf{R}_{pca}))^{-1} \quad (31)$$

where $\text{diag}(\mathbf{R}_{pca})$ is the matrix containing the diagonal elements of $\mathbf{R}_{pca}$.

It is clear that the transformed model $\mathbf{R}_t$ has the form of a collection of regression models. Now, because of the "normalization" of the residuals, the predictive ability of $\mathbf{R}_t$ can be compared directly to a collection of PLS models $\mathbf{R}_{pls}$ as proposed in Eqs. (11) to (14).

The system used in Wise and co-workers (1990), will again be used as an example to compare the predictive ability of the methods. For the test, data was generated exactly as described above, except that the noise level was varied from 0 to 1.0 times the noise level specified previously in increments of 0.1. In each case 1000 samples were generated and PCA and PLS models, $\mathbf{R}_t$ and $\mathbf{R}_{pls}$, were formed. In each case the PCA model retained 5 PCs. The number of latent variables in each of the PLS models was optimized based on prediction error using cross validation. The predictive ability of the models was then tested on a new data set with the same noise level as the calibration set. Note that the only difference in each of the calibration and test data sets was the noise magnitude. Identical input and noise sequences were used.

The results of a subset of the prediction error tests are shown in Fig. 5. The results are shown in terms of the percentage decrease in total sum of squares prediction error for the PLS model as compared to the PCA model, *e.g.*, for variable 7 and a noise level of 0.9 times the base noise level, the PLS model sum of squared error was almost 60% less than for the PCA model. Note how the difference in predictive ability of the models increases as the noise level is increased. It is also evident that there is a larger difference for some variables than for others.

The results of this experiment show why PLS based residuals might be superior to those based on PCA. With the PCA model transformed in this manner, the change in a residual given a change in a variable due to error is identical for both the PCA and PLS models. However, the PLS model residuals under

normal conditions are substantially smaller than the PCA residuals. Therefore, for the PLS model the change due to error is relatively larger, and should be more easily detected.

Comparison of the models $\mathbf{R}_t$ and $\mathbf{R}_{pls}$ themselves showed some interesting trends. When the noise level is zero the models are identical, as might be expected. As the noise level is increased, $\mathbf{R}_t$ changes very little. This is consistent with experience which has shown that it is possible to identify an accurate PCA model even when the noise level is quite high. $\mathbf{R}_{pls}$, on the other hand, changes a great deal as the noise becomes very high. For optimum predictive ability, the model requires fewer latent variables, as would be expected. In general, the coefficients of the model tend to get smaller, though for some variables they may get larger. This trend is indicative of the shift towards fewer latent variables, which tends to spread the predictive ability of the model over more variables rather than concentrating it on a few.

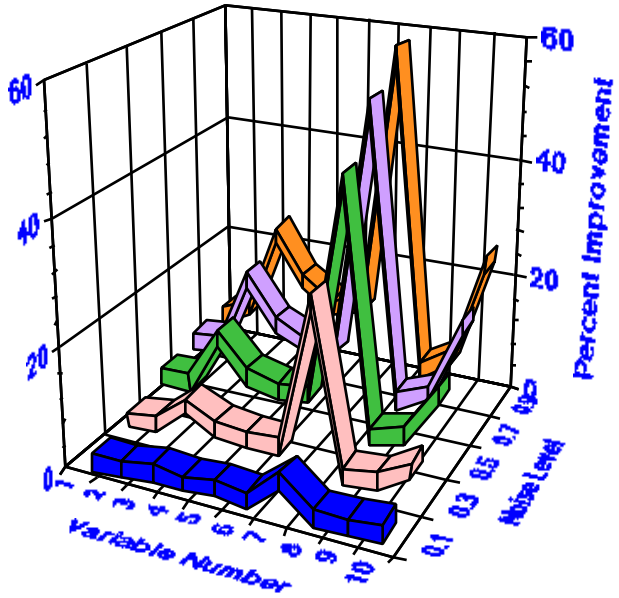

Fig. 5. Percent improvement in sum of squared prediction error for PLS models over PCA models.

While it can be expected that the PLS model $\mathbf{R}_{pls}$ will be more sensitive to changes in the process data, it cannot be expected that the residuals will behave as in the PCA case. In general, PLS models do not produce zero-mean residuals. Furthermore, to the extent that $\mathbf{R}_{pls}$ lies outside of the subspace spanned by $\mathbf{R}_{pca}$, it can be expected that some state information will be mapped into the residuals. Because the states are usually autocorrelated, the PLS residuals will be also. This is shown in Fig. 6, where the autocorrelation function (ACF) of the PLS residuals is shown along with the ACF of the process outputs. For this test the process was driven by a Pseudo Random Binary Sequence (PRBS). It is apparent that there is a good deal of autocorrelation in the residuals. There is no autocorrelation in the corresponding PCA residuals, which are not shown. If the autocorrelation in the states changes (perhaps due to a change in input behavior or a disturbance), then the autocorrelation of the

PLS residuals would also be affected. Any change in the autocorrelation of the residuals would have the affect of making the control limits invalid. Thus, any PLS detection scheme based on a particular correlation structure in the process states would become invalid given a change in the state behavior.

In spite of the autocorrelation in the residuals, the detection limits for the PLS models (adjusted to account for the autocorrelation) are better than for the PCA models. This is shown in Fig. 7, where the approximate detection limits are given for the PCA and PLS models. Here the limits are based on a 20 sample window and 99% confidence. The PCA limits are calculated directly from theory as shown in Wise (1990). The PLS limits are based on the observed mean and standard deviation of 20 sample subgroups of the process data. Note that the PLS limits are better for all variables, but the difference is more significant for some than others.

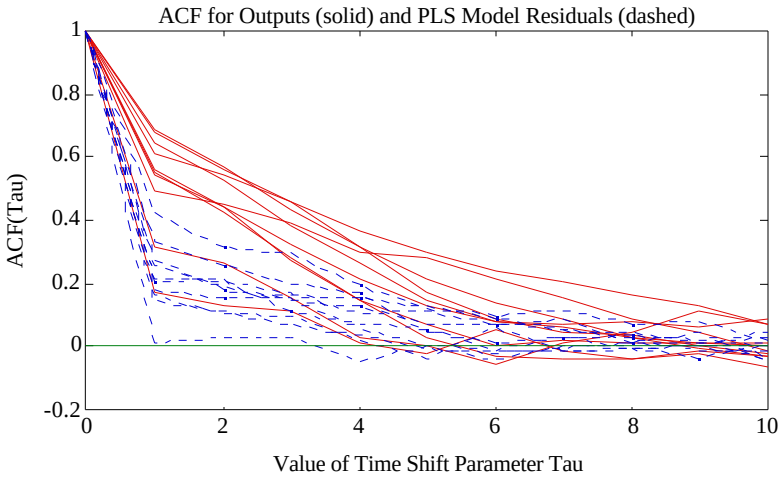

Fig. 6. Autocorrelation function of process outputs and PLS residuals from data set with PRBS input.

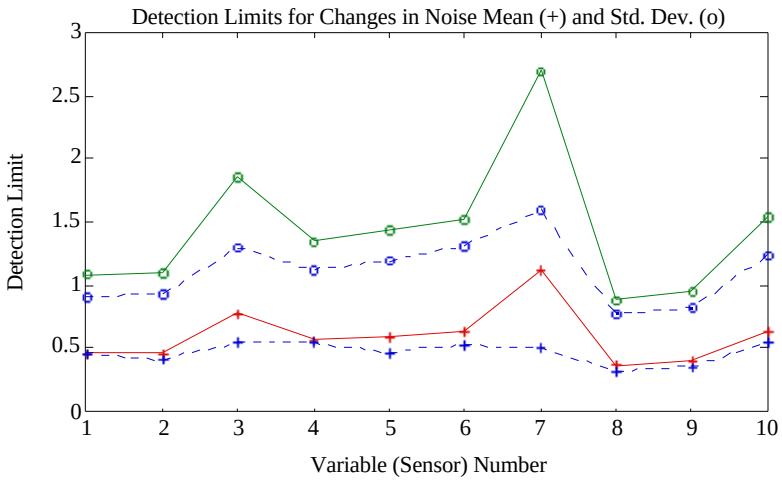

Fig. 7. Detection limits for changes in measurement noise mean and standard deviation for PCA (solid lines) and PLS models (dashed lines).

It was pointed out in Wise and co-workers (1990) that the assumption that any process has a finite number of states is clearly an approximation, although generally a useful one. Thus, it is an approximation to say that the data from any process (where measurements are made at different locations) is intrinsically rank deficient. The PLS monitoring method, on the other hand, makes no assumptions about the intrinsic rank of the process data and, therefore, about the order of the process producing it. Instead, the PLS models are built up individually and the criteria for the models is the predictive residual error. Thus there is no "cutoff" approximation as in the PCA models. It could be said, however, that in each PLS model the number of factors used for prediction is an estimate of the number of process states that are relevant in the prediction of each output.

Comparison of PCA and PLS Monitoring

The ability of the PCA and PLS models to accurately identify sensor bias and was tested through simulation. Two new 1000 sample data sets were generated using the example model. In the first case the model was driven by white noise (as in the calibration set) and in the second case the model was driven by a PRBS, which has considerably more power at low frequencies. The data sets were broken into 50 segments of 20 points each Bias was then added to each variable in the segment in turn, and the residuals were calculated and tested for significance. The "failed sensors" were identified based on the ratio of the residual mean or standard deviation to the appropriate limits, *i.e.*, the output with the largest relative residual was assumed bad.

The results of the error detection simulation are shown in Table 1. The table is divided into two sections corresponding to the case of white noise input to the model or the PRBS sequence which is autocorrelated. The results for the PCA model are the four columns on the left while the PLS model results are the four columns on the right. Each of the four columns corresponds to a different level of added bias or noise. The basis for the bias and noise levels is the scaled outputs, *i.e.*, a bias of 0.5 indicates that a bias of 0.5 units was added to the scaled variable. The original scaling, which resulted in a mean zero unit variance calibration set, was used in the test. A noise error of 1.0 corresponds to adding white noise of unit standard deviation to the scaled outputs. The row labeled "Correct" indicates the number of times the method correctly identified the proper variable has having added bias. "None" indicates the number of times there were no variables over the limits. "Incorrect" indicates the number of times an out of bounds variables was detected but the wrong variable was indicated as faulty. The sum of the "Incorrect" and "None" categories is given in the final line of the table.

TABLE 1. Error Detection Simulation Results for Bias Errors

Model Input White Noise

	PCA Model				PLS Model			
Bias Size	0.5	1.0	1.5	2.0	0.5	1.0	1.5	2.0
Correct	141	423	485	499	229	485	499	500
None	286	29	1	0	187	4	0	0
Incorrect	73	48	14	1	84	11	1	0
Co. + Inco.	359	77	15	1	271	15	1	0

Model Input Correlated PRBS

	PCA Model				PLS Model			
Bias Size	0.5	1.0	1.5	2.0	0.5	1.0	1.5	2.0
Correct	153	438	492	500	186	412	473	492
None	298	24	1	0	89	2	0	0
Incorrect	49	38	7	0	225	86	27	8
Co. + Inco.	347	62	8	0	314	88	27	8

It is clear that, under the same conditions as the calibration data, (in this case a white noise input), the PLS model performance is superior. When the model input is changed to PRBS, however, the PLS model performance suffers considerably while the PCA model performance is essentially unchanged. This degradation of the PLS model performance is due to the mapping of state information into the residuals. When the autocorrelation in the states changes, the PLS residual limits become invalid.

Improving PLS Model Robustness

Like the PCA models, the unmodified PLS models will become useless after one sensor failure. However, outputs from failed sensors can be estimated based on the collection of PLS models in the same way as for PCA models. It is possible to solve for the values of $\mathbf{x}_b$ which minimizes the magnitude of the residual vector. The lack of symmetry in the $\mathbf{R}_{pls}$ matrix causes a minor complication, however, the result is

$$\mathbf{x}_b = -\mathbf{x}_g((\mathbf{R}_{21} + \mathbf{R}_{12}^T)/2)\mathbf{R}_{11}^{-1} \tag{32}$$

where it is understood that the $\mathbf{R}_{pls}$ matrix has been partitioned as in Eq. (18). It is interesting to note that, while the residuals on the bad variables becomes zero in the PCA case, it does not in the PLS case. Instead, the new residuals on the bad variables are:

$$\mathbf{r}_b = -\mathbf{x}_g\mathbf{R}_{ave} + \mathbf{x}_g\mathbf{R}_{21} \tag{33}$$

where

$$\mathbf{R}_{ave} = (\mathbf{R}_{21} + \mathbf{R}_{12}^T)/2 \tag{34}$$

So to the extent that $\mathbf{R}_{21}$ is not equal to $\mathbf{R}_{12}^T$ the residuals on the fixed variables will not be zero.

CONCLUSIONS

This article has shown how the robustness and sensitivity of MSPC can be improved. Replacing a failed sensor output with its PCA-based estimate allows for detection of additional failures, which would otherwise be obscured. Furthermore, this method produces residuals that are nearly identical to the residuals from an entirely new PCA model.

PLS can be used in manner similar to PCA for failure detection. It has been shown that the use of predictive ability as the basis for model selection makes PLS more sensitive than PCA. The major drawback of PLS is that it produces residuals that are autocorrelated, due to mapping of state information into the residuals. The autocorrelation in the residuals can change when the autocorrelation in the states changes, invalidating the calculated limits.

REFERENCES

Åstrom, K. J., and B. Wittenmark (1984). Computer Controlled Systems. Prentice-Hall, New Jersey.

Geladi, P., and B. R. Kowalski (1986). PLS tutorial. Anal. Chim. Acta., 185, 1-17.

Jackson, J. E. (1981). Principal components and factor analysis: part 1-principal components, J. Qual. Tech., 12, 201-213.

Kresta, J., J. F. MacGregor, and T. E. Marlin (1990). Multivariate statistical monitoring of process operating performance, Submitted to Can. J. Chem. Eng., In Press.

Kwakernaak, H., and R. Sivan (1972). Linear Optimal Control Systems. Wiley-Interscience, New York.

Lorber, A., L. E. Wangen, and B. R. Kowalski (1987). A theoretical foundation for the PLS algorithm. J. Chemometrics., 1, 19-31.

MacGregor, J. F. (1989). Multivariate statistical methods for monitoring large data sets from chemical processes. AIChE Annual Meeting, San Francisco.

Malinowski, E. R. (1977). Determination of the number of factors and the experimental error in a data matrix. Anal. Chem., 49, 612-617.

Malinowski, E. .R. (1987). Theory of the distribution of error eigenvalues resulting from principal component analysis with applications to spectroscopic data. J. Chemometrics, 1, 33-40.

Sage, A. P., and C. C. White III (1977). Optimum Systems Control, 2nd Ed., Prentice-Hall, New Jersey.

Strang, G. (1980). Linear Algebra and Its Applications. Academic Press, New York.

Veltkamp, D. J., B. R. Kowalski, N. L. Ricker, and B. M. Wise (1990). Multivariate statistical process control using principal component analysis", Submitted to J. Chemometrics.

Wise, B. M. (1991). Adapting Multivariate Analysis for Monitoring and Modeling of Dynamic Processes. Ph.D. Dissertation, University of Washington, Seattle.

Wise, B. M., and N. L. Ricker (1989). Feedback strategies in multiple sensor systems. AIChE Symposium Series, Vol. 85, No. 267, pp. 19-23.

Wise, B. M. D. J. Veltkamp, N. L. Ricker, B. R. Kowalski, S. M. Barnes, and V. Arakali (1991). Application of multivariate statistical process control (MSPC) to the West Valley slurry-red ceramic melter process. Waste Management '91 Proceedings, Tucson, In Press.

Wise, B. M., N. L. Ricker, and D. J. Veltkamp (1989). Upset and sensor failure detection in multivariate processes. AIChE Annual Meeting, San Francisco.

Wise, B. M., N. L. Ricker, D. J. Veltkamp, and B. R. Kowalski (1990). A theoretical basis for the use of principal component models for monitoring multivariate processes. Process Control and Quality, 1, 41-51.