Isolation Enhanced Principal Component Analysis

Janos Gertler and Weihua Li
School of Information Technology, George Mason University, Fairfax, VA 22030

Yunbing Huang and Thomas McAvoy
Chemical Engineering Dept., University of Maryland, College Park, MD 20742

Principal component analysis (PCA) may reduce the dimensionality of plant models significantly by exposing linear dependences among the variables. While PCA is a popular tool in detecting faults in complex plants, it offers little support in its original form for fault isolation. However, by utilizing the equivalence between PCA and parity relations, all the powerful concepts of analytical redundancy may be transferred to PCA. Following this path, it is shown how structured residuals, which have the same isolation properties as analytical redundancy residuals, are obtained by PCA. The existence conditions of such residuals are demonstrated, as well as how disturbance decoupling is implied in the method. The effect of the presence of control constraints in the training data is analyzed. Statistical testing methods for structured PCA residuals are also outlined. The theoretical findings are fully supported by simulation studies performed on the Tennessee Eastman process.

Introduction

Multivariate statistical methods have recently become the tools of choice in the monitoring of complex chemical processes (MacGregor, 1989; Kresta et al., 1991; Piovoso et al., 1992a,b; Wise et al., 1995). Principal component analysis (PCA) is used to model normal process behavior, and unusual events (faults) are then detected by referencing the observed behavior against this model. The great advantage of PCA is that it may significantly reduce the dimension of the plant model. However, it provides little support for fault isolation. Contribution charts (MacGregor et al., 1994) and multiblock approaches (Chen and McAvoy, 1997) have been proposed recently, but they offer only limited solutions.

As we have shown recently, a close equivalence exists between PCA and parity relations (Gertler and McAvoy, 1997). Parity relations belong to the analytical redundancy (AR) methods (Chow and Willsky, 1984), which rely on an explicit model of the plant. AR methods have well developed fault isolation capabilities which utilize the special (structured or directional) design of the residual set (Gertler and Singer, 1990; Gertler, 1991). By virtue of the equivalence, these concepts may be transferred into the PCA framework. The resulting method combines the convenience of statistical-type modeling with the powerful isolation properties of AR.

In our previous article, we outlined the generation of structured residuals by a set of ''partial'' PC models. Here we will describe a method which does this algebraically, relying on a single PC model. First, we will give an introduction to parity relations. This will be followed by a brief review of the standard PCA technique, leading to the main result of the article. We will also address the existence conditions of structured PCA residuals. It will then be shown that disturbance decoupling is included in the approach automatically, subject to the limitations well-known from AR studies. The main results, originally derived for sensor and actuator faults, will then be extended to plant faults. This will be followed by an analysis of the effect and handling of control constraints. Finally, the background and some strategies of testing the structured residuals will be outlined.

Results of simulation studies performed on the Tennessee Eastman process (Downs and Vogel, 1993) are presented. The results of these studies have been in complete agreement with the theory.

Note that, in order to focus on the fundamental concepts, we assume throughout this article that the system is static and linear. The parity relation technique is well developed

Correspondence concerning this article should be addressed to J. Gertler.

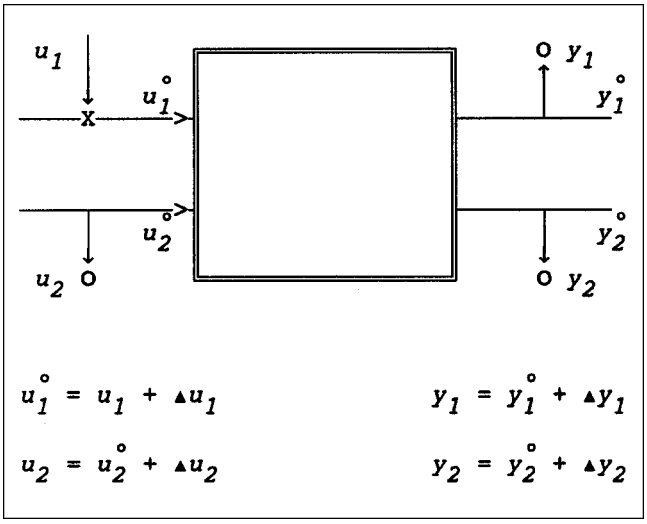

Figure 1. Variables and faults.

for linear dynamic systems (Chow and Willsky, 1984; Gertler and Singer, 1990), and applications to nonlinear problems have also been reported (Gertler et al., 1995). While PCA is fundamentally a linear technique, several extensions to nonlinear systems are known in the literature (Etezadi-Amoli and McDonald, 1983; Dong and McAvoy, 1996) and some dynamic extensions have also been developed (Chen and McAvoy, 1997). The generalization of the results described in this article to nonlinear systems is the subject of current research, and some early results will be reported in (Huang et al., 1999).

Background

System description

Consider a linear static system with observed (manipulated or measured) inputs $u(t) = [u_1(t), \ldots, u_k(t)]'$ and observed (measured) outputs $y(t) = [y_1(t), \ldots, y_m(t)]'$. These inputs and outputs are subject to actuator and sensor faults $\blacktriangle u(t)$ and $\blacktriangle y(t)$. Reserving the symbols $u(t)$ and $y(t)$ for the *observed* values of the variables, these are related to the *true* values (those acting on or arising from the plant) $u^0(t)$ and $y^0(t)$ as

$$u(t) = u^0(t) \pm \blacktriangle u(t) \qquad y(t) = y^0(t) + \blacktriangle y(t) \quad (1)$$

Figure 1 illustrates the variables and faults for a system with two inputs (one manipulated and one measured) and two outputs (both measured).

We assume that the variables are centered relative to their mean and that each variable and its fault are normalized with the variable's standard deviation. We make three additional assumptions here which will be removed later in the article:

(a) No disturbances act in the system;
(b) There are no plant-faults;
(c) No control constraints act on the variables.

Under these assumptions, the true plant variables are linked by the linear relationship

$$y^0(t) = Au^0(t) \quad (2)$$

By combining u and y into a single vector

$$x(t) = [y'(t) \quad u'(t)]' \quad (3)$$

one may write Eqs. 1 and 2 as

$$x(t) = x^0(t) + \blacktriangle x(t) \quad (4)$$

$$Bx^0(t) = \mathbf{0} \quad \text{where} \quad B = [I \quad -A] \quad (5)$$

Structured parity relations

In the analytical redundancy approach, the model matrix A is known (either from first principles or from prior identification of the plant). Parity relations (Gertler and Singer, 1990) rely on Eq. 5, applied to the observed variables. With Eqs. 4 and 5

$$Bx(t) = B \blacktriangle x(t) = o(t) \quad (6)$$

where $o(t)$ is the set of *primary residuals* which are computed from the observations but depend only on the faults. Further residuals are obtained by the transformation

$$r(t) = Wo(t) \quad (7)$$

Structured Residuals. Structured residuals are so designed that each residual is sensitive to a particular subset of faults. With a suitable combination of residual structures, each fault triggers a characteristic response, also called its *fault code*. In the structure or incidence matrix, the rows are the residual structures while the columns are the fault codes. An *I* element means that the concerned residual does respond to the concerned fault while an *O* means it does not. A structure is ''isolating'' if each column is different and ''strongly isolating'' if the columns stay different from valid columns when *I*s are turned into *O*s. Strong isolation is best achieved by *column canonical* structures, in which each fault code has the same number of *O*s, each in a different pattern. The following is a column canonical structure for four faults $\blacktriangle x_1 \ldots \blacktriangle x_4$ and four residuals $r_1 \ldots r_4$.

	$\blacktriangle x_1$	$\blacktriangle x_2$	$\blacktriangle x_3$	$\blacktriangle x_4$
r_1	*O*	*I*	*I*	*O*
r_2	*O*	*O*	*I*	*I*
r_3	*I*	*O*	*O*	*I*
r_4	*I*	*I*	*O*	*O*

Note that column canonical structures have received considerable attention in coding theory where they are called *m out of n codes* or constant weight codes (Rao and Fujiwara, 1989). Note also that such structures are normally designed for single faults; if multiple faults occur, the resulting fault codes exhibit an increased number of *I*s (in many cases, they are all *I*s). Thus, multiple faults cannot usually be isolated from fault codes designed for single faults, though their presence is possible to detect. With extended sets of residuals, however, strongly isolating (column canonical) structures may also be designed for multiple faults (such as double or triple) (Gertler, 1998).

When structured residuals are generated by Eq. 7, then the ith row of W, w_i' is so chosen that the zeros assigned for the ith row of the structure matrix be implemented, that is, the ith residual be decoupled from all faults which appear with an O in that row. This requires that

$$w_i' B^i = \mathbf{0} \tag{8}$$

where B^i contains those columns of the B matrix which belong to the faults assigned for zero response in the ith residual structure. To make the transformation completely defined, an additional nonzero condition is necessary, which may be the desired value of one of the nonzero fault-responses (Gertler and Singer, 1990; Gertler, 1998).

Existence Conditions. If there are m primary residuals, then any w_i' vector contains m elements. It follows then from Eq. 7 that the number of zeros in any row cannot exceed $m-1$. More precisely the matrix B^i has to satisfy the rank restriction

$$\text{Rank } B^i \le m-1 \tag{9}$$

Further, to assure that there are no unwanted zeros in the row structure, that is, the residual is not decoupled also from faults which appear with I in the concerned row, the condition

$$\text{Rank}\begin{bmatrix} B^i & b_g \end{bmatrix} = \text{Rank } B^i + 1 \tag{10}$$

must be satisfied for all columns b_g of B outside B^i. If the conditions 9 and 10 are not met, the selected row-structure is not *attainable*. Finally, if any two columns of B are linearly dependent, then it is not possible to construct a residual set which distinguishes between the two faults; these faults are not *isolable*.

Standard PCA

Consider the training data set $x(\tau)$, $\tau = 1, \ldots, N$, where it is assumed that no fault is present, that is, $x(\tau) = x^0(\tau)$. Perform a singular value decomposition for

$$\mathbf{\Omega} = [\, x(1), \ldots, x(N)\,] \tag{11}$$

as

$$\mathbf{\Omega} = \sum_{j=1}^{k+m} \sigma_j^* q_j \gamma_j' \tag{12}$$

where q_j and γ_j are the eigenvectors of $\mathbf{\Omega}\mathbf{\Omega}'$ and $\mathbf{\Omega}'\mathbf{\Omega}$, respectively, and σ_j^* are the singular values (the non-negative square-roots of the eigenvalues of $\mathbf{\Omega}\mathbf{\Omega}'$) (Kresta et al., 1991). In the presence of m linear relations among the variables, $\text{Rank}\,\mathbf{\Omega} = k$; then m of the singular values are zero, so the upper limit of the summation effectively becomes k

$$\mathbf{\Omega} = \sum_{j=1}^{k} \sigma_j^* q_j \gamma_j' \tag{13}$$

Note that when working with real physical data, precise zero singular values are seldom encountered, due to the effects of noise, disturbances, and plant nonlinearities. These issues will be addressed later in this article; here we assume that some of the singular values are exactly zero.

It follows from Eqs. 11 and 13 that any column of $\mathbf{\Omega}$ can be expressed with the eigenvectors $q_1 \ \ldots \ q_k$ (those which belong to nonzero singular values) as

$$x(\tau) = \sum_{j=1}^{k} \sigma_j^* q_j \gamma_{j\tau}' = \sum_{j=1}^{k} q_j s_j(\tau) \tag{14}$$

Here the eigenvectors $q_1 \ \ldots \ q_k$ are the *principal directions* and the coefficients $s_j(\tau)$ are the *scores*. The latter are the orthogonal projections of the observations on the eigenvectors which can also be computed as

$$s_j(\tau) = q_j' x(\tau) \qquad j = 1 \ \ldots \ k \tag{15}$$

Equations 14 and 15 may also be written in compact form as

$$x(\tau) = \tilde{Q}\, s(\tau) \tag{16}$$

$$s(\tau) = \tilde{Q}'\, x(\tau) \tag{17}$$

where

$$\tilde{Q} = [\, q_1 \ \ldots \ q_k\,] \tag{18}$$

$$s(\tau) = [\, s_1(\tau) \ \ldots \ s_k(\tau)\,]' \tag{19}$$

Note that in order to characterize the data as in Eqs. 14 and 16, the full singular value decomposition (Eq. 12) of $\mathbf{\Omega}$ is not really necessary. What is needed are the eigenvectors q_j and the singular values. These are most conveniently obtained from the empirical covariance matrix

$$R = \frac{1}{N}\mathbf{\Omega}\mathbf{\Omega}' = \frac{1}{N}\sum_{\tau=1}^{N} x(\tau)\, x'(\tau) \tag{20}$$

The covariance matrix R has the same eigenvectors as $\mathbf{\Omega}\mathbf{\Omega}'$ and its eigenvalues σ_j^2 are related to those of $\mathbf{\Omega}\mathbf{\Omega}'$ as $\sigma_j^2 = (\sigma_j^*)^2/N$.

When the above PC model is applied to an observation $x(t)$ outside the training set, this in general is not perfectly described by the expansion in the principal directions $q_1 \ \ldots \ q_k$ but an error (residual) may arise

$$x(t) - \tilde{Q}\, s(t) = x(t) - \tilde{Q}\,\tilde{Q}'\, x(t) = \boldsymbol{\epsilon}(t) \ne \mathbf{0} \tag{21}$$

Structured Residuals by PCA

Primary residuals

We focus our attention now on the residual $\boldsymbol{\epsilon}(t)$. First we will show that

$$\boldsymbol{\epsilon}(t) = Q\, Q'\, x(t) \tag{22}$$

where

$$Q = [\, q_{k+1} \ \cdots \ q_{k+m} \,] \tag{23}$$

contains the eigenvectors which belong to zero singular values. To see this, consider

$$\mathbf{\Gamma} = [\, \tilde{Q} \quad Q \,] = [\, q_1 \ \cdots \ q_{k+m} \,] \tag{24}$$

Since the eigenvectors are orthonormal, $\mathbf{\Gamma}'\mathbf{\Gamma} = I$, from which $\mathbf{\Gamma}' = \mathbf{\Gamma}^{-1}$. Therefore, $\mathbf{\Gamma}\mathbf{\Gamma}' = \mathbf{\Gamma}\mathbf{\Gamma}^{-1} = I$. Thus from Eq. 21

$$\boldsymbol{\epsilon}(t) = \mathbf{\Gamma}\mathbf{\Gamma}' x(t) - \tilde{Q}\,\tilde{Q}'\, x(t)$$

$$= [\, \tilde{Q} \quad Q \,] \begin{bmatrix} \tilde{Q}' \\ Q' \end{bmatrix} x(t) - \tilde{Q}\,\tilde{Q}'\, x(t) = Q\,Q'\, x(t)$$

We may now reinterpret the PC model as follows. The eigenvectors $q_1 \ \cdots \ q_k$ span the subspace where the training set exists. Since the training set is assumed to be fault-free, this is also where any future fault-free variable-vector may exist. We will refer to this as the *representation subspace*. The remaining eigenvectors $q_{k+1} \ \cdots \ q_{k+m}$ span the orthogonal complement of the representation subspace; this is where the residuals exist. This will be referred to as the *residual subspace*; it is m-dimensional. Note that any true variable is orthogonal to the residual subspace thus

$$Q' x^0(t) = \mathbf{0} \tag{25}$$

Note that in Eq. 22 the residuals are represented as linear combinations of the eigenvectors $q_{k+1} \ \cdots \ q_{k+m}$ (those which belong to zero singular values). This approach was used by Lou et al. (1986) in connection with the application of singular value decomposition to approximate disturbance decoupling (see also Gertler, 1998).

Equation 22 describes the residuals in terms of the original x coordinates. It will be more advantageous to express them in the residual space, in terms of the $q_{k+1} \ \cdots \ q_{k+m}$ coordinates as

$$e(t) = Q' x(t) = Q' \blacktriangle x(t) \tag{26}$$

where we took into account Eqs. 4 and 25. Clearly, the transformation between the two representations is $\boldsymbol{\epsilon}(t) = Qe(t)$. We will consider Eq. 26 the *primary PC residual*; it is computed from the observation $x(t)$ but depends only on the faults $\blacktriangle\, x(t)$.

Relationship between Q' and B

Equation 26 shows a remarkable similarity to Eq. 6. This is not simply formal; one can show that

$$Q' = MB \tag{27}$$

where M is a full-rank square matrix. To see this, observe that both Q' and B contain m rows which are $k + m$ vectors. The rows of Q' are orthonormal eigenvectors; they are orthogonal to the eigenvectors $q_1 \ \cdots \ q_k$ spanning the representation subspace. The rows of B are not orthonormal but they are linearly independent, by virtue of the unit matrix B contains (see Eq. 5). By Eq. 5, they are orthogonal to any possible fault-free variable vector $x^0(t)$; thus, they are also orthogonal to the representation subspace. This implies that both sets of vectors span the residual subspace, from which Eq. 27 follows.

Denote by q_{*j}, $j = 1 \ \cdots \ m + k$, the columns of the Q' matrix (these are m vectors). Then Eq. 27 applies to any column as

$$q_{*j} = Mb_j \qquad j = 1 \ \cdots \ m + k \tag{28}$$

This implies that any linear dependences among the columns of B transfer to the columns of Q' and vice versa. That is, Q' and B have the same column-rank properties. Further, applying Eq. 28 to the first m columns, and observing that $[\, b_1 \ \cdots \ b_m \,] = I$ (see Eq. 5), it follows that

$$[\, q_{*1} \ \cdots \ q_{*m} \,] = M[\, b_1 \ \cdots \ b_m \,] = M \tag{29}$$

Note that the training data determines the residual subspace but, with multiple zero singular values, the eigenvectors spanning this subspace are only partially defined. Thus, any actual Q' contains arbitrary choices. Therefore, the transformation M cannot be determined using only the knowledge of the explicit model B. In turn, the model B determines the representation subspace, but not the exact position of the eigenvectors which span it. The latter depend also on how much activity the training data contains along the various input variables.

Structured residual design

Let us expand the primary residual (Eq. 26) as

$$e(t) = [\, q_{*1} \ q_{*2} \ \cdots \ q_{*(m+k)} \,] \begin{bmatrix} \blacktriangle\, x_1(t) \\ \blacktriangle\, x_2(t) \\ \vdots \\ \blacktriangle\, x_{m+k}(t) \end{bmatrix} \tag{30}$$

This equation indicates that the response to a fault $\blacktriangle\, x_j(t)$ is

$$e\big(t | \blacktriangle\, x_j\big) = q_{*j} \blacktriangle\, x_j(t) \tag{31}$$

Thus the column q_{*j} can be considered as the *image* of the jth fault in the residual subspace.

New residuals may be obtained from the primary PC residuals by the transformation

$$p(t) = Ve(t) \tag{32}$$

The rows of the transformation $\boldsymbol{\nu}_i'$ are of the same dimension as the fault images. However, the transformed residual $p(t)$ may be of any dimension. The decouple an element $p_i(t)$

of the residual from a fault $\blacktriangle x_j$, the row-transformation $\boldsymbol{\nu}_i'$ has to be orthogonal to the image q_{*j}. That is

$$\boldsymbol{\nu}_i' q_{*j} = 0 \tag{33}$$

Further, to decouple a residual-element from several faults

$$\boldsymbol{\nu}_i' (Q')^i = \mathbf{0} \tag{34}$$

is required, where $(Q')^i$ contains the images of the faults assigned for decoupling from the ith residual (zeros in the structure matrix). Equation 34 is the fundamental design equation for structured PCA residuals; notice its similarity to Eq. 8.

To supplement the homogeneous decoupling conditions (Eq. 34), the length of the $\boldsymbol{\nu}_i'$ vector may be chosen as unity. Then, $p_i(t)$ is the orthogonal projection of the primary residual $e(t)$ on $\boldsymbol{\nu}_i'$. Clearly, for the decoupled faults, the projection of $e(t|\blacktriangle x_j)$ is zero.

Existence conditions follow from the column properties of the Q' matrix. For the transformation (Eq. 34) to exist

$$\text{Rank}(Q')^i \le m-1 \tag{35}$$

is required; this usually means that the number of zeros per row must not exceed $m-1$. To avoid unintended zeros

$$\text{Rank}\left[(Q')^i \quad q_{*g}\right] = \text{Rank}(Q')^i + 1 \tag{36}$$

must be satisfied for any column q_{*g} outside $(Q')^i$. If Eq. 36 is not met, the selected row-structure is not attainable. Finally, if any two columns of Q' are linearly dependent, the respective faults are not isolable.

Observe that Eqs. 35 and 36 are formally identical with Eqs. 9 and 10. Further, by virtue of Eq. 28, the structural restrictions (nonattainability and nonisolability) following from the PCA model are identical with those that would arise from the explicit model of the same system.

Design procedure

The design consists of the following steps:

(1) Perform a standard PCA and obtain the Q' matrix;

(2) Determine m and design a structure matrix;

(3) Check Q' for rank-defects and modify the structure if necessary;

(4) Implement the row-transformations by Eq. 34.

Note that the distinction between zero and nonzero singular values may not be quite clear-cut; we will comment on this in the next section.

We wish to emphasize the absolute parallelism between the PCA-based procedure and the parity relation design. In fact, one may even switch to parity relations, by computing the explicit model from the PCA model, using Eqs. 27 and 29.

Disturbance Decoupling

Disturbances are unmeasured plant inputs which are normally present in the data. They affect the outputs and thus interfere with fault detection and isolation. Whenever possible, the diagnostic algorithm is to be so designed that it is insensitive to the disturbances.

In the presence of disturbances $d(t) = [d_1(t) \ \ldots \ d_l(t)]'$, the input-output relationships (Eqs. 2 and 5) become

$$y^0(t) = Au^0(t) + Fd(t) \tag{37}$$

$$Bx^0(t) = Fd(t) \tag{38}$$

It is assumed that the scalar time-sequences $d_1(t)$, $d_2(t)$, ..., $d_l(t)$, $t = 1, 2, \ldots$, are linearly independent. (If this is not the case, the disturbance-set is reduced by combining the linearly dependent elements.)

Parity relations

In the analytical redundancy framework, F is assumed to be known, either from the physical understanding of the process or from experiments. (This, of course, may not always be a realistic assumption.) Insensitivity to the disturbances is achieved by decoupling from them the residuals. This amounts to including the disturbances in the incidence matrix with full zero columns. The transformation (Eq. 8) is then obtained as

$$w_i'[F \quad B^i] = \mathbf{0} \tag{39}$$

An equivalent alternative procedure is to first obtain a set of intermediate residuals which are decoupled from all the disturbances, and then apply selective fault-decoupling to these in order to achieve structured fault responses.

The implementation condition (Eq. 9) is now

$$\text{Rank}[F \quad B^i] \le m-1 \tag{40}$$

Clearly, the presence of F reduces by $\text{Rank}\, F \le l$ the number of faults from which a particular residual can be decoupled. Equation 40 also implies that the maximum number of disturbances, with independent columns in F, from which decoupling is possible without losing fault detection is $m-1$ while that without losing isolation is $m-2$. The attainability condition (Eq. 10) becomes

$$\text{Rank}\left[F \quad B^i \ b_g\right] = \text{Rank}[F \quad B^i] + 1 \tag{41}$$

Thus, the presence of the F matrix increases the possibility of nonattainability and nonisolability.

PCA

In the PCA framework, we assume that each element of the disturbance vector is present in the training data. Note that the size of the various components is irrelevant. Decoupling from disturbances which are missing from the training data, but show up in the monitoring phase, is not possible.

The presence of disturbances in the training data increases by $\text{Rank}\, F$ the dimension of the representation subspace and, thus, results in the reduction of the dimension of the residual subspace accordingly. To see this, consider Eq. 38; with an

appropriate transformation T the disturbances can be eliminated from the equation

$$TBx^0(t) = TFd(t) = \mathbf{0} \tag{42}$$

This clearly requires

$$TF = \mathbf{0} \tag{43}$$

The rows of T have m elements and need to be orthogonal to $\text{Rank}\,F$ independent vectors, thus

$$\text{Rank}\,T = m - \text{Rank}\,F \geq m - l \tag{44}$$

Since $\text{Rank}\,B = m$, the number of independent relations among the variables, with the disturbances eliminated, is

$$\text{Rank}[T \quad B] \leq \text{Rank}\,T = m - \text{Rank}\,F \tag{45}$$

This is also the dimension of the residual subspace. Note that the above transformation is equivalent to using some of the original equations to express the disturbances and substituting them into the remaining equations.

With the disturbances present in the training data, PCA will automatically generate the reduced residual subspace. This can then be used to design residuals structured for the faults. (This is like first designing parity relation residuals decoupled only from the disturbances.) Disturbances represented in the training data will not show up in those residuals. The design restrictions are the same as with parity relations. In particular:

- If $\text{Rank}\,F = m$ (it cannot be greater since F has m rows), then the dimension of the residual subspace is zero, and there is no fault detection;
- If $\text{Rank}\,F = m - 1$, then the dimension of the residual subspace is normally one, faults can be detected but not isolated;
- The residuals become insensitive to any fault whose column b_j is colinear with F, see Eq. 41;
- In general, the reduced residual subspace increases the possibility of nonattainability and nonisolability.

As we mentioned before, the distinction between zero and nonzero singular values may not be obvious in practice. This is because model errors (nonlinearities) and noise also appear as disturbances and their combined number may exceed $m - 1$. Therefore, the residuals need to be threshold tested. We will comment on threshold selection in a later section. Here, we only mention that, in a practical situation, the dimension of the residual subspace is chosen somewhat arbitrarily, considering the trade-off between more design freedom if the dimension is higher and greater fault sensitivity if it is lower.

Plant Faults

Plant faults are unmeasured inputs which act distinctly from the observed inputs and outputs. They are normally not present in the data; when they appear, we wish the diagnostic algorithm to detect and isolate them.

Mathematically, plant faults appear the same way as disturbances (see Eqs. 37 and 38). In the presence of plant faults $f(t) = [f_1(t) \ \ldots \ f_h(t)]$, the input-output relationships (Eqs. 2 and 5) become

$$y^0(t) = Au^0(t) + Ef(t) \tag{46}$$

$$Bx^0(t) = Ef(t) \tag{47}$$

It is assumed that the scalar time-sequences $f_1(t)$, $f_2(t)$, ..., $f_h(t)$, $t = 1, 2, \ldots$, are linearly independent.

Parity relations

In the analytical redundancy framework, the gain matrix E is assumed to be known, either from the physical understanding of the process or from experiments. (Such information, of course, may be difficult to obtain.) Then the plant faults are included in the structured residual design, by expanding the structure matrix to cover them as well, and appending E to the B matrix. The computation of the residuals, of course, is still done with the original B matrix.

PCA

In the PCA framework, plant faults can be properly detected and isolated if it is possible to obtain training data, not only in the fault-free situations, but also with plant faults present, one at a time. Clearly, this requirement may be quite difficult to satisfy in practice (note that it is not necessary to know the size of the plant faults). Then the following procedure may be followed.

(1) First the residual subspace (Q matrix) is established from fault-free training data;

(2) Then residuals are computed from the training data containing one plant fault at a time, as

$$e(\tau | f_j) = Q' x(\tau | f_j) \qquad \tau = 1 \ \ldots \ N_j, \qquad j = 1 \ \ldots \ h \tag{48}$$

(3) Next the fault images q_{*m+k+j}, $j = 1 \ \ldots \ h$, are established for each plant fault. This can be done by applying PC analysis to the residual covariance matrices

$$R_{ej} = \frac{1}{N_j} \sum_{\tau=1}^{N_j} e(\tau | f_j) e'(\tau | f_j) \qquad j = 1 \ \ldots \ h \tag{49}$$

The dominant eigenvector for each plant fault will be the respective fault image.

(4) Finally, the plant-fault images are appended to the Q' matrix and the plant faults included in the structure matrix, then structured residuals are designed for the expanded set.

Of course, in the monitoring phase the residuals are still computed with the original Q' matrix.

Note that if plant faults appear in the monitoring phase for which the residuals have not been designed, the observed fault-codes will not agree with any of the valid codes, and in many cases will contain all Is. Thus, in general, such faults may be detected as unexplained faults, but may not be isolated.

Control Constraints

Recall that the equations describing plant behavior are valid for the *true values* of the variables. If applied to the observed values, these equations return nonzero residuals in the presence of faults, disturbances, noise, and so on. This is the basis for residual generation.

There is another group of relationships, imposed by control actions, which act on the *observed* (measured and commanded) *values*. Such control actions include:

(1) *Feedback control* keeping variables constant. Since feedback is based on sensory measurement, this constraint applies to the measured value of the controlled variable, rather than its true value. Note that with the variables being centered relative to their mean, the only possible constant (centered) value is zero.

(2) *Ratio control* keeping variables linearly related. This has two possible implementations:

(i) One input is actuated to be proportional to the measured value of another input, in which case the constraint applies to the measured value and the actuator command; and

(ii) Two inputs are actuated to a preset ratio (without measurement), in which case the constraint applies to the actuator commands.

Control constraints always return zero residuals, no matter what faults, and so on, are present, except when the control action breaks down completely.

The set of control constraints is described as

$$Cx(t) = \mathbf{0} \tag{50}$$

A constant variable is included as a special case with only one nonzero element in the respective row of C. Equation 50, combined with Eq. 6, leads to the extended residual equation

$$\begin{bmatrix} B \\ C \end{bmatrix} x(t) = \begin{bmatrix} B \\ \mathbf{0} \end{bmatrix} \blacktriangle x(t) = \begin{bmatrix} o(t) \\ \mathbf{0} \end{bmatrix} \tag{51}$$

Parity relations

In analytical redundancy type design, where the distinction between plant and control equations is obvious, the latter are simply left out, so that the residuals are designed and generated entirely based upon Eq. 6.

Note that control constraints may create some difficulties in the analytical redundancy framework if the plant model is obtained by system identification. Data collected under control constraints is generally poorly exciting and leads to uncertainties in the identified parameters. These, in turn, may cause problems in the isolation of the concerned sensor and actuator faults. More about this can be found in Gertler (1998).

PCA

In the PCA framework, the presence of control constraints in the training data increases the dimension of the residual subspace the same way as plant equations do. Once the training data have been collected under such constraints, we do not know of any way to eliminate or distinguish them in the PCA model. If the increased residual space is used for structured design, as if all the equations had originated from the plant, then imperfect decoupling of the concerned sensor and actuator faults will ensue.

To see this, consider the extended system model (Eq. 51). Assume that training data are collected with the constraints present; PCA then yields an extended residual subspace Q'_C. This is related to the extended system model as

$$Q'_C = M_C \begin{bmatrix} B \\ C \end{bmatrix} = [\, M_{CI} \quad M_{CII} \,] \begin{bmatrix} B \\ C \end{bmatrix} = M_{CI} B + M_{CII} C \tag{52}$$

where M_C is a full-rank square transformation and M_{CI} and M_{CII} are appropriate submatrices.

For any column of Q'_C, Eq. 52 implies

$$q_{C*j} = M_{CI} b_j + M_{CII} c_j \tag{53}$$

The primary residual is computed as

$$e(t) = Q'_C x(t) = M_{CI} B x(t) + M_{CII} C x(t) = M_{CI} B \blacktriangle x(t) \tag{54}$$

where we utilized Eq. 51. This implies that the response to the jth fault is

$$e\big(t|\blacktriangle x_j\big) = M_{CI} b_j \blacktriangle x_j \tag{55}$$

Now if we design a residual $p_i(t)$, with the intention of decoupling from $\blacktriangle x_j$, based on the column q_{C*j}, so that $\boldsymbol{\nu}'_i q_{C*j} = 0$, then the actual residual in response to $\blacktriangle x_j$ will be

$$\begin{aligned} p_i\big(t|\blacktriangle x_j\big) &= \boldsymbol{\nu}'_i e\big(t|\blacktriangle x_j\big) = \boldsymbol{\nu}'_i M_{CI} b_j \blacktriangle x_j \\ &= \boldsymbol{\nu}'_i \Big[\underbrace{M_{CI} b_j + M_{CII} c_j}_{q_{C*j}} - M_{CII} c_j \Big] \blacktriangle x_j = -\boldsymbol{\nu}'_i M_{CII} c_j \blacktriangle x_j \end{aligned} \tag{56}$$

In general, this is nonzero if $c_j \neq 0$.

To avoid this kind of problem, the control constraints must be removed from the PC model. To achieve this, feedback loops have to be operated with (at least two) different setpoints and ratio controllers with different ratios when collecting training data.

Note that control constraints are part of a more general problem of sufficient excitation during the training of a PC model; this is the subject of ongoing research. Further issues of interest are the detection of the presence of control constraints from the PC model, and the interplay between control constraints and disturbances.

Statistical Testing

If PCA is applied to practical data, exact zero singular values are never obtained. This is due to noise and disturbances always present in the data. The latter include the effects of

nonlinearities and dynamics which, under linear static modeling, act as input-dependent disturbances. The lack of zero singular values has two consequences:

(1) An arbitrary decision has to be made as to where to draw the line between true plant variables and faults, that is, which eigenpairs to consider part of the representation subspace and which of the residual subspace.

(2) The residuals need to be threshold tested for fault detection and diagnosis. The thresholds are primarily based on the variance information provided by PCA, which may be supplemented with assumptions concerning the statistical properties of the residuals.

In this section, we address the issues of testing strategy and threshold selection.

Residual variances

The way the primary residuals $e(t)$ are computed from the observations $x(t)$, see Eq. 26, amounts to a ''diagonalizing'' transformation, leading to a residual vector whose elements are uncorrelated. To see this, consider the empirical covariance matrix of the primary residuals over the training data

$$R_e = \frac{1}{N}\sum_{\tau=1}^{N} e(\tau)\, e'(\tau) \tag{57}$$

With Eq. 26, this becomes

$$R_e = \frac{1}{N}\sum_{\tau=1}^{N} Q'x(\tau)\, x'(\tau)Q = Q'RQ \tag{58}$$

where R is the covariance matrix of the data, see Eq. 20. This latter can be expressed with its eigenstructure as

$$R = \mathbf{\Gamma} S^2 \mathbf{\Gamma}' \tag{59}$$

where $\mathbf{\Gamma} = [\tilde{Q}\ Q]$, see Eq. 24, and S is the diagonal matrix of the singular values. Now $Q'\mathbf{\Gamma} = [\mathbf{0}\ I]$, so that

$$R_e = Q'\mathbf{\Gamma} S^2 \mathbf{\Gamma}' Q = [\mathbf{0} \quad I]\begin{bmatrix} S_I^2 & \mathbf{0} \\ \mathbf{0} & S_{II}^2 \end{bmatrix}\begin{bmatrix} \mathbf{0} \\ I \end{bmatrix} = S_{II}^2 \tag{60}$$

Here S_I and S_{II} are $k \cdot k$ and $m \cdot m$ diagonal submatrices of S, associated with the representation and residual subspaces, respectively. Thus, the elements $e_1(t)\ \ldots\ e_m(t)$ of the primary residual vector $e(t)$ are uncorrelated and their variances are the squared singular values $\sigma_{k+1}^2\ \ldots\ \sigma_{k+m}^2$.

By virtue of the transformation (Eq. 32), the elements of the transformed residual vector $p(t)$ are correlated. Their covariance matrix is obtained as

$$R_p = \frac{1}{N}\sum_{\tau=1}^{N} p(\tau)\, p'(\tau) = VS_{II}^2 V' \tag{61}$$

Variances of the individual residual elements $p_i(t)$ are in the main diagonal of R_p; they can also be computed as

$$\sigma_{pi}^2 = \boldsymbol{\nu}_i' S_{II}^2 \boldsymbol{\nu}_i = \sum_{j=1}^{m} \nu_{ij}^2 \sigma_{k+j}^2 \tag{62}$$

where $\boldsymbol{\nu}_i'$ is the ith row transformation and $\boldsymbol{\nu}_i$ is its transpose, with ν_{ij} being its jth element.

Testing strategies

The theoretical design of tests requires the knowledge of the distribution of the quantity to be tested, in our case the residuals. While we do have information on some parameters of the fault-free distribution (zero mean and variances as computed above), the type of the distribution is unknown. Usually, the assumption of normal distribution is introduced, which may be reasonable for the noise-induced part of the residual but more questionable as the disturbances and model errors are concerned.

In connection with structured residuals, three testing strategies are possible:

(1) *Individual Testing of Scalar Residuals.* This is the strategy that leads to the binary fault codes usually associated with structured residuals. The individual test thresholds are proportional to the standard deviations computed from the singular values in Eq. 62. With the usual assumption of normal distribution, the threshold can be designed for a specified false alarm rate.

(2) χ^2 *Detection Testing.* Under the assumption of normal distribution, the statistic

$$\xi(t) = \sum_{j=1}^{m} e_j^2(t)/\sigma_{k+j}^2 \tag{63}$$

obeys χ^2 distribution with m degrees of freedom. This statistic expresses the normalized length of the primary residual vector. Comparing it against χ^2 thresholds offers an easy indication of the presence of any fault.

(3) *Constrained GLR Testing.* This is a more sophisticated testing strategy which utilizes the full information in the residual covariance matrix to isolate faults from structured residuals. Conditional estimates of the residual means are computed with the various fault hypotheses, under the geometric constraints posed by the respective structures, and likelihood functions obtained with those mean are compared. This strategy also requires the assumption of normal distribution. More on this approach can be found in Gertler (1998).

Each of the above tests may be performed on a single observation or on the average of a sequence of observations. Tests based on averages are, of course, more powerful (in terms of the probability of missed detections under a set probability of false alarms). If the tests are conducted on averages, then the variances on which the thresholds rest need to be reduced accordingly. For an uncorrelated sequence of observations, which is the case usually assumed in static systems, this simply implies dividing the variances by the number of observations in the sequence. As for the more general problem of testing correlated sequences, see Gertler (1998).

Simulation Results

Tennessee Eastman Process

This process is a simulation of a real plant that has been disguised for proprietary reasons (Downs and Vogel, 1993). The process produces two products G and H from four reac-

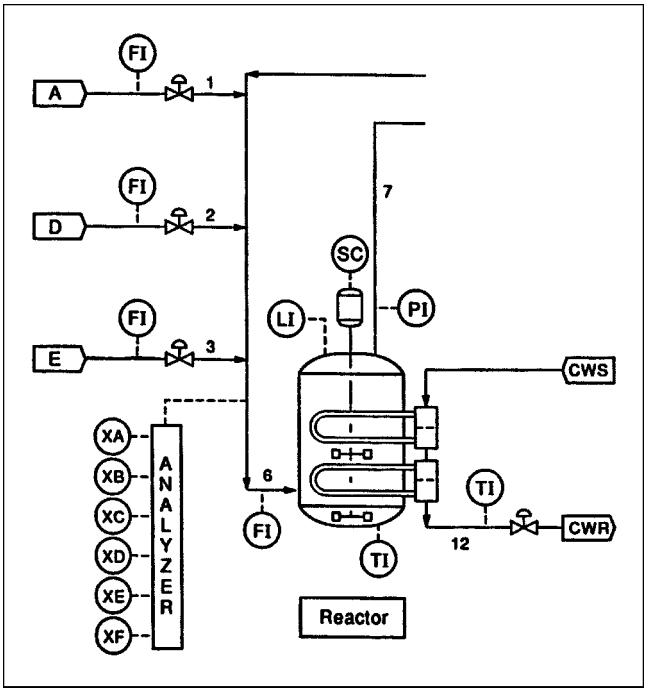

Figure 2. Tennessee Eastman reactor.

tants *A*, *C*, *D*, and *E*. Also present are an inert *B* and a byproduct *F*. The process has five major units: a reactor, a product condenser, a vapor/liquid separator, a recycle compressor, and a product stripper. The gaseous reactants are fed to the reactor where they react to form liquid products. The gas-phase reactions are catalyzed by a nonvolatile catalyst dissolved in the liquid phase. The products leave the reactor as vapors along with unreacted feeds and the catalyst remains in the reactor. The process has 41 measurements and 12 manipulated variables. The modes of process operation are set by the G/H mass ratios. There are six modes of process operation at three different G/H mass ratios.

To simplify the demonstration, we only considered the reactor unit (Figure 2). The reactor has 12 variables:

temperature T;	feed mole-fractions $X_A \ldots X_F$;
pressure P;	coolant temperature T_c;
level L;	coolant valve position V_c; and
feed flow F	

Of these T, P, L, and T_c are measured outputs, F and $X_A \ldots X_F$ are measured inputs, and V_c is a manipulated input. Of the six mole-fractions, at most, five are independent.

The complete system is simulated with nine of the control loops operating with constant setpoints. Three variables T, P, and L are controlled with their setpoints varied. Of these, the temperature is controlled by manipulating the coolant valve in a cascade loop where the reference for the coolant temperature is varied inside the loop. Because of the nine fixed setpoints, only three of the inputs to the reactor are independent, including the coolant valve. Thus, the flow and one of the mole fractions determine the remaining mole fractions. By means of varying the three setpoints, the additional control constraints were removed.

Table 1. PCA Results: Scenario 1

PC No.	Eigenvalue	% Variance	% Var. Total
1	4.10	58.60	58.60
2	1.87	26.76	85.37
3	1.02	14.61	99.97
4	1.57×10^{-3}	0.03	100.00
5	1.83×10^{-4}	0.00	100.00
6	1.96×10^{-5}	0.00	100.00
7	3.77×10^{-8}	0.00	100.00

Table 2. Residual Subspace: Scenario 1

$$Q'=\begin{bmatrix} -0.7907 & -0.1788 & 0.1614 & -0.1506 & -0.2240 & 0.4526 & -0.1978 \\ 0.3349 & -0.5392 & 0.0574 & 0.2729 & 0.1718 & 0.2403 & -0.6573 \\ -0.0710 & -0.3998 & -0.0048 & -0.5297 & -0.2265 & -0.6701 & -0.2327 \\ 0.0036 & -0.0203 & -0.0027 & 0.5954 & -0.7797 & -0.1924 & -0.0087 \end{bmatrix}$$

The investigations were performed in three scenarios. In each one, the PC model was trained with a data set of 100 samples. Then, the diagnostic algorithm was designed and tested with data containing one fault at a time. In each scenario, the algorithm was designed for six faults, those associated with F, P, L, T, T_c, and V_c. Throughout the simulation, low-level noise was added to all outputs.

Scenario 1

X_A is included as a measured input. For the seven variables, the results yielded by PCA are shown in Table 1.

The representation space has dimension 3, the residual space dimension 4. The Q' matrix is shown in Table 2.

A possible strongly isolating structure is given in Table 3.

The transformation V and the fault-transfer VQ' are shown in Table 4

Clearly, the first six columns of VQ' follow the patterns of the structure matrix. The last column of VQ' belongs to X_A,

Table 3. Strongly Isolating Structure: Scenario 1

	F	P	L	T	T_c	V_c
r_1	O	I	I	O	O	I
r_2	I	O	I	I	O	O
r_3	I	I	O	O	I	O
r_4	O	O	O	I	I	I

Table 4. Transformation and Fault-Transfer: Scenario 1

$$V=\begin{bmatrix} -0.3391 & -0.8742 & -0.3476 & 0.0057 \\ -0.6616 & 0.4920 & -0.3885 & 0.4113 \\ 0.3113 & -0.8776 & -0.1932 & 0.3091 \\ 0.0126 & 0.0055 & -0.0637 & 0.9979 \end{bmatrix}$$

$$VQ'=\begin{bmatrix} 0 & 0.6709 & -0.1032 & 0 & 0 & -0.1317 & 0.7225 \\ 0.7170 & 0 & -0.0778 & 0.6847 & 0 & 0 & -0.1057 \\ -0.5252 & 0.4886 & 0 & 0 & -0.4178 & 0 & 0.5576 \\ 0 & 0 & 0 & 0.6275 & -0.7655 & -0.1423 & 0 \end{bmatrix}$$

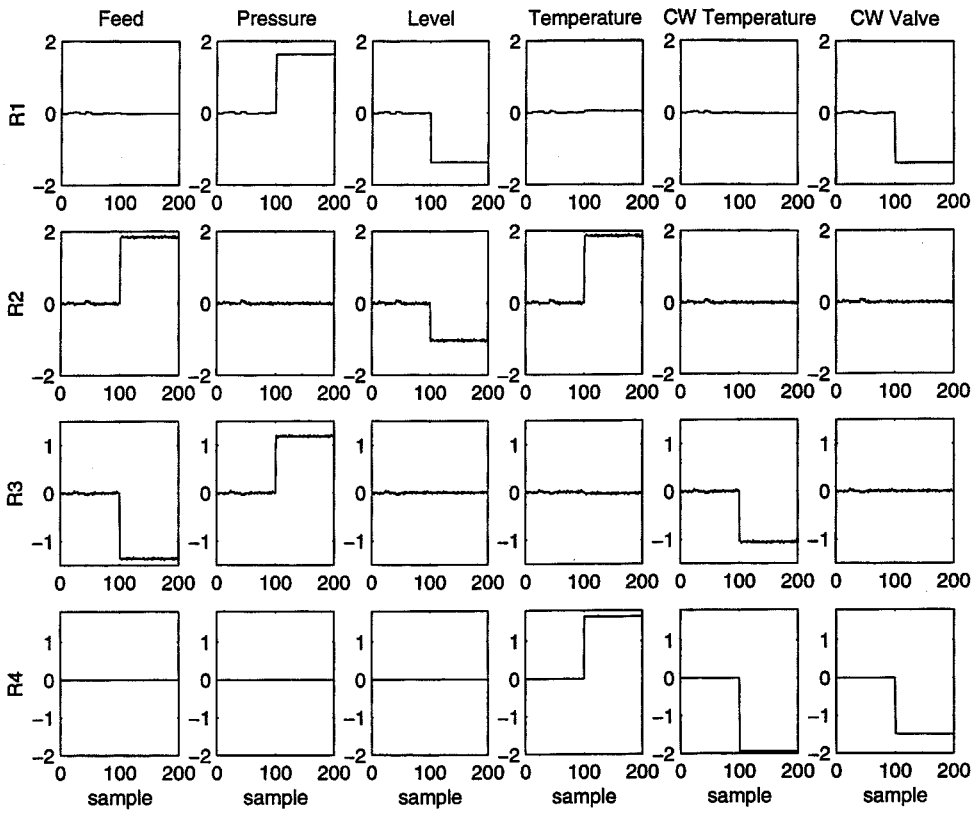

Figure 3. Fault responses in Scenario 1.

a variable whose accompanying fault was not included in the design of the structure matrix. As can be seen, the response to this fault is not a valid code under the column-canonical design, in that it has three *I*s (instead of two). As such, it is an indication of an ''unidentified'' fault (which it really is) or a multiple fault (two or three of the first three faults simultaneously).

The fault responses are plotted in Figure 3. The plots are arranged in accordance with the structure matrix; each row shows the response of the same residual to the various faults while each column is the response of the residual set to a particular fault. Clearly, in each fault situation, some of the residuals do respond while the others do not. The fault codes obtained are exactly what is expected, that is, the simulated behavior of the diagnostic system is in full agreement with the structured design.

Table 5. PCA Results: Scenario 2

PC No.	Eigenvalue	% Var.	% Var. Total
1	3.94	65.68	65.68
2	1.04	17.38	83.06
3	1.01	16.91	99.97
4	1.52×10^{-3}	0.03	100.00
5	2.01×10^{-6}	0.00	100.00
6	1.77×10^{-11}	0.00	100.00

Table 6. Residual Subspace: Scenario 2

$$Q'=\begin{bmatrix} -0.8164 & 0.0297 & 0.1321 & -0.1099 & -0.2066 & 0.5103 \\ 0.2551 & 0.2477 & 0.0067 & 0.3675 & 0.5221 & 0.6826 \\ -0.0838 & -0.0687 & 0.0002 & -0.7552 & 0.6458 & -0.0312 \end{bmatrix}$$

Scenario 2

Now X_A is not taken into account as an input, thus it acts as a disturbance. The PCA results are given in Table 5.

The representation space is still dimension 3, but the residual space has been reduced to dimension 3. The Q' matrix for the six variables is given in Table 6.

Now the submatrix formed of the first 3 columns has a rank-defect. This limits the attainable structures, but it is still possible to design a strongly isolating structure, as shown in Table 7.

Table 7. Strongly Isolating Structure: Scenario 2

	F	P	L	T	T_c	V_c
r_1	*O*	*I*	*I*	*O*	*I*	*I*
r_2	*O*	*I*	*I*	*I*	*O*	*I*
r_3	*I*	*O*	*I*	*O*	*I*	*I*
r_4	*I*	*O*	*I*	*I*	*I*	*O*
r_5	*I*	*I*	*O*	*I*	*O*	*I*
r_6	*I*	*I*	*O*	*I*	*I*	*O*

Table 8. PCA Results: Scenario 3

PC No.	Eigenvalue	% Var.	% Var. Total
1	3.31	55.13	55.13
2	1.08	18.07	73.21
3	1.00	16.67	89.88
4	6.07×10^{-1}	10.12	100.00
5	1.29×10^{-4}	0.00	100.00
6	1.43×10^{-7}	0.00	100.00

The fault-transfer matrix and the response plots, not shown in this article, again follow the structure.

Scenario 3

Now an additional disturbance is applied upstream to the feed. The PCA result is given in Table 8.

The residual space is further reduced to dimension 2. The Q' matrix is given in Table 9.

The first 3 columns are now pairwise linearly dependent; thus, the first three faults are not isolable. A possible structure under this constraint is shown in Table 10.

In summary, the simulation results have been found to be in complete agreement with the theory.

Conclusion

Utilizing the equivalence between principal component analysis and parity relations, we have demonstrated the generation of structured residuals in the PCA framework. These residuals have the same powerful isolation properties as analytical redundancy residuals.

Based on experimental data, the eigenstructure of the variable covariance matrix is obtained. The variable space is divided into representation and residual subspaces according to the singular value magnitudes. The primary residuals are the projections of the observations to the residual subspace. Subjecting the primary residuals to an additional transformation, in accordance with a predesigned structure matrix, residuals selectively sensitive to the various sensor and actuator faults are obtained.

The existence conditions of structured residuals are discussed, and it is shown how disturbance decoupling is implied in the method. The design is extended to plant faults, provided experimental data with such faults present is available. The effect of the presence of control constraints in the training data is addressed. Finally, the background and methods of the statistical testing of structured PCA residuals are outlined. The theoretical findings are fully supported by simulation studies performed on the Tennessee Eastman process.

While the various concepts have clear geometric interpretations, the described procedure hinges on an algebraic transformation. There is an alternative way of achieving the same result, a procedure we termed *partial PCA* (Gertler and McAvoy, 1997). In the partial PCA design, structured residuals are obtained by performing separate eigenstructure analyses for subsets of the variables, chosen in accordance with the structure matrix. In a linear system, these residuals have exactly the same properties as the ones obtained algebraically. Partial PCA is equivalent to a technique used with parity relations where transformed models are identified directly from experimental data in structures defined by the structure matrix.

Table 9. Residual Subspace: Scenario 3

$$Q'=\begin{bmatrix} 0.0512 & -0.3275 & -0.0657 & -0.4696 & -0.2051 & -0.7894 \\ -0.0008 & 0.0048 & 0.0010 & -0.6183 & 0.7682 & 0.1660 \end{bmatrix}$$

Table 10. Partially Isolating Structure: Scenario 3

	F	P	L	T	T_c	V_c
r_1	O	O	O	I	I	I
r_2	I	I	I	O	I	I
r_3	I	I	I	I	O	I
r_4	I	I	I	I	I	O

The extension of the described methods to discrete linear dynamic systems is straightforward. In accordance with the fact that such systems are modeled by autoregressive-moving average discrete transfer functions, a number of past values (as many as the dynamic order of the system) of each input and output variable need to be included in the PCA. With this, the size of the algorithm increases (and the clear geometric interpretation is lost), but otherwise the procedures remain unchanged.

Extension to nonlinear systems is more difficult, especially when relying on algebraic transformation. However, in the partial PCA framework, where no transformation is involved, nonlinear representations (such as polynomial models or neuronets) may be introduced. Once the structure matrix has been established, partial PCA's can be performed, combined with such nonlinear representations. This area, which would correspond to parity relations with directly identified nonlinear models (Gertler, 1998), is a subject of ongoing research.

Acknowledgment

The research reported here has been supported by NSF under grant no. CTS-9714891.

An earlier version of this article was presented at the IFAC Workshop on On-Line Fault Detection and Supervision in the Chemical Process Industries, Lyon, France, June 4–5, 1998. The first author would like to acknowledge a conversation with Mark Kramer in June 1995 when he suggested that the link between PCA and parity relations might be worth looking into.

Literature Cited

Chen, G., and T. J. McAvoy, "Multi-Block Predictive Monitoring of Continuous Processes," *IFAC ADCHEM Symposium*, Banff, Canada, 73 (1997).

Chow, E. Y., and A. S. Willsky, "Analytical Redundancy and the Design of Robust Failure Detection Systems," *IEEE Trans. on Automatic Control*, **29**, 603 (1984).

Dong, D., and T. J. McAvoy, "Nonlinear Principal Component Analysis Based on Principal Curves and Neural Networks," *Comput. and Chem. Eng.*, **20**, 65 (1996).

Downs, J. J., and E. F. Vogel, "A Plant-Wide Industrial Process Control Problem," *Comp. Chem. Eng.*, **17**, 245 (1993).

Etezadi-Amoli, J., and R. P. McDonald, "A Second Generation Nonlinear Factor Analysis," *Psychometrica*, **48**, 315 (1983).

Gertler, J., and D. Singer, "A New Structural Framework for Parity Equation Based Fault Detection and Isolation," *Automatica*, **26**, 381 (1990).

Gertler, J., ''Analytical Redundancy Methods in FDI—Survey and Synthesis,'' *IFAC SAFEPROCESS Symp*., Baden Baden, Germany, Vol. 1, 9 (1991).
Gertler, J., M. Costin, X. Fang, Z. Kowalczuk, M. Kunwer, and R. Monajemy, ''Model-Based Diagnosis for Automotive Engines—Algorithm Development and Testing on a Production Vehicle,'' *IEEE Trans. on Control Systems Technology*, **3**, 61 (1995).
Gertler, J., and T. J. McAvoy, ''Principal Component Analysis and Parity Relations—A Strong Duality,'' *Preprints of IFAC SAFEPROCESS Symp*., Hull, UK, Vol. 2, 837 (1997).
Gertler, J., *Fault Detection and Diagnosis in Engineering Systems*, Marcel Dekker, New York (1998).
Huang, Y., J. Gertler, and T. J. McAvoy, ''Fault Isolation by Partial PCA and Partial NLPCA,'' *14th IFAC World Congress*, Beijing, China, in press (1999).
Kresta, J. V., J. F. MacGregor, and T. E. Marlin, ''Multivariate Statistical Monitoring of Processes,'' *Can. J. Chem. Eng*., **69**, 35 (1991).
Lou, X. C., A. S. Willsky, and G. C. Verghese, ''Optimally Robust Redundancy Relations for Failure Detection in Uncertain Systems,'' *Automatica*, **22**, 333 (1986).
MacGregor, J., ''Multivariate Statistical Methods for Monitoring Large Datasets from Chemical Processes,'' AIChE Meeting, San Francisco (1989).
MacGregor, J., C. Jaeckle, C. Kiparissides, and M. Koutodi, ''Process Monitoring and Diagnosis by Multiblock Methods,'' *AIChE J*., **40**, 826 (1994).
Piovoso, M. J., K. A. Kosanovich, and P. K. Pearson, ''Monitoring Process Performance in Real Time,'' *Amer. Control Conf*., Chicago, p. 2359 (1992a).
Piovoso, M. J., K. A. Kosanovich, and J. P. Yuk, ''Process Data Chemometrica,'' *IEEE Tr. Instr. and Meas*., **41**, 262 (1992b).
Rao, T. R. N., and E. Fujiwara, *Error-Control Coding for Computer Systems*, Prentice Hall, Englewood Cliffs, NJ (1989).
Wise, B. M., N. B. Gallagher, and J. F. MacGregor, ''The Process Chemometrics Approach to Process Monitoring and Fault Detection,'' *Proc. IFAC Workshop on On-Line Fault Detection and Supervision in the Chemical Process Industries*, Newcastle, UK, 1 (1995).

Manuscript received July 2, 1998, and revision received Nov. 12, 1998.