Sensor validation for smart structures

Michael I. Friswell[*a] and Daniel J. Inman[b]

[a] Department of Mechanical Engineering, University of Wales Swansea, Swansea SA2 8PP, UK
[b] Center for Intelligent Material Systems and Structures, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061, USA

ABSTRACT

Structures with a large number of embedded sensors are becoming more common, and this refined spatial information can be used to advantage in damage location and model validation. These sensors could be accelerometers, strain gauges, piezoceramic patches, PVDF film sensors, or optical fibre sensors. This approach requires that the sensors are functioning correctly, which on a smart structure operating in the field should be continuous and automatically monitored. This paper considers possible approaches to sensor validation, based on the assumption that a model of the structure is available. The aim is to make use of the natural data redundancy since there will more sensors than modes in the data. The validation approaches considered are based on hypothesis testing based on a number of techniques, such as modal filtering. The methods are demonstrated on simple examples that exercise their strengths and weaknesses.

Keywords: Sensor Validation, Smart Structures, Modal Filtering, Principal Component Analysis

1. INTRODUCTION

Smart structures have the potential to place a large number of distributed actuators and sensors on a structure. The correct functioning of active control and health monitoring requires that the sensors be functioning. Errors introduced by faulty sensors can cause undamaged areas to be identified as damaged, or control system performance to be inadequate, or certainly less than optimal. In many civil structures applications for health monitoring (such as bridges), ambient loads must be used for excitation. These loads are not known and may be measured or estimated as part of the health monitoring algorithm, which requires a large number of sensors.

Sensor validation, where the sensors are confirmed to be functioning during operation, seems to have received little attention in the smart structures community. The critical aspect in smart structures is that there are usually more sensors than excited modes. This redundancy may be used, together with a modal model of the structure, to validate the sensor functionality.

The control and chemical engineering community have considered the sensor validation problem, and have used models and sensor redundancy to good effect [1-9]. However, these approaches usually use the faulty sensor to predict the response and look for errors between predictions and measurement. Clearly using the faulty sensor in the prediction process will propagate errors to the predicted responses. Often neural networks, or artificial intelligence approaches are used for the analysis. This paper is based on the assumption that only the lower modes of the structure are usually excited, producing a large redundancy in the data. This has similarities to the principal component analysis used in chemical plant [10-12]. The approach is demonstrated on a beam structure, although the method is completely general and may be applied to any structure for which the modal model is available.

* Correspondence: E-mail: m.i.friswell@swansea.ac.uk; Tel: +44 (0)1792 295217; Fax: +44 (0)1792 295676

In *Fifth European Conference on Smart Structures and Materials*, Pierre François Gobin, Clifford M. Friend, Editors, Proceedings of SPIE Vol. 4073 (2000) • 0277-786X/00/$15.00

2. RESPONSE TO ARBITRARY FORCING

The standard equations of motion for the response of structure to an arbitrary forcing, $\mathbf{f}(t)$, are

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{K}\mathbf{x} = \mathbf{f}(t) \tag{1}$$

where $\mathbf{x}$ is the structure's response and $\mathbf{M}$ and $\mathbf{K}$ are the mass and stiffness matrices. Damping will be considered later. By applying the modal transformation, based on the mass normalised mode shapes $\{\phi_i\}$,

$$\mathbf{x} = \Phi\mathbf{q} \tag{2}$$

where $\mathbf{q}$ is the modal participation factor and Φ is the mode shape matrix, $\Phi = [\,\phi_1\ \phi_2 \ldots \phi_n\,]$. We get,

$$\ddot{q}_i + \lambda_i q_i = Q_i \tag{3}$$

where λ_i is the i th eigenvalue, and $Q_i = \phi_i^T \mathbf{f}$ is the i th modal force. Note that Equation (2) implies the response of the structure is a sum of the responses in each mode. Explicitly this is,

$$\mathbf{x} = \sum_{i=1}^{n} q_i \phi_i \tag{4}$$

The summation is given over all of the n modes of the structure. Typically only the lower modes are important because the force excitation is concentrated in these modes. Also, the anti-aliasing filters required for the data sampling mean that the response at the higher modes will be filtered out. This constraint that the response lies in a subspace generated by the lower modes of the structure is central to the proposed sensor validation techniques. The sensor validation could be performed directly, by computing the response orthogonal to the modal subspace, however weighting should be included to ensure an orthogonal basis to the subspace. Furthermore, the probability that a sensor is damaged may be computed by the approach proposed here, and this probability may be utilised in supervisory fault diagnosis systems.

In practice not all the degrees of freedom of the system in equation (1) will be measured. Let $\mathbf{H}$ denote the measurement matrix, then the measured outputs, $\mathbf{y}$, are

$$\mathbf{y} = \mathbf{H}\mathbf{x} \tag{5}$$

Of course the number of sensors must be greater than the number of modes of interest to enable this data redundancy to exist. In smart structures, where a large number of sensors are distributed throughout a structure, this condition is likely to be met.

Thus far the model has been assumed accurate and the damping neglected. Undamped mode shapes are usually estimated accurately from the model. An alternative would be to use the measured mode shapes (assuming all the sensors are functioning correctly). In any case, faults in the sensors will only be identified if they causes changes in the response of a greater magnitude than the errors in the estimated mode shape. Damping has not been considered, but if the damping ratios of the system are small then the response will be contained (approximately) in the subspace defined by the undamped (and therefore real) modes. In any case, the modes with low damping, having approximately real modes, will be excited most strongly.

3. SENSOR VALIDATION CONCEPTS

Although there is redundancy in the data, based on the number of sensors and the number of modes excited, it is still not straight-forward to identify those sensors that are damaged. When all sensors are working it is possible to estimate the modal contributions to the response and therefore produce a predicted response that will give some idea of the accuracy of the model of structure and the extent of measurement noise. However if a sensor is damaged, then using data from this sensor to estimate the modal participation factors will propagate the errors from the faulty channel through the estimate of modal response to the estimate of the response in all channels. Thus to predict faulty sensors the sensors are split into two groups. If S represents the set of all sensors then these two groups are,

$$S_f = \{\text{sensors assumed to be faulty}\}$$
$$S_w = \{\text{sensors assumed to be working}\} \tag{6}$$

Note that these sets are disjoint so that

$$S_f \cap S_w = \{\ \} \qquad S_f \cup S_w = S \tag{7}$$

Note that the distribution of faulty and working sensors seems to have been determined at the outset. Of course in practice we do not know which sensors will be faulty and so we have to try every potential subset of faulty sensors. This approach has parallels with the subset selection techniques in parameter identification [13, 14], and for sensor validation, as for parameter estimation, the difficulty is to determine which parameter subset is optimal. Note that for sensor validation, the number of assumed working sensors should be at least as great as the number of modes of interest. For each set we can test 2 hypotheses, namely,

- that using S_w the assumed model fits the data
- that at least some sensors in S_f are faulty

(8)

What we do is take the outputs from the sensors in S_w, and use the model to estimate the modal participation factors. If the predictions based on this model accurately reproduce the data then we have confidence that both the model and the sensors are operational. The estimated modal participation factors and the model may then be used to predict the outputs at the candidate faulty sensor locations. Faulty sensors are then determined by errors in these predictions.

4. VALIDATION VIA MODAL FILTERING

Central to the proposed strategy for sensor validation is a modal model of the structure and also the estimation of the modal participation factors during operation. At any time instant, t_k, the measured output is

$$\mathbf{y}(t_k) = \mathbf{y}_k = \mathbf{H}\,\boldsymbol{\Phi}\,\mathbf{q}(t_k) = \mathbf{H}\,\boldsymbol{\Phi}_r\,\mathbf{q}_r(t_k) + \mathbf{H}\,\boldsymbol{\Phi}_d\,\mathbf{q}_d(t_k) \tag{9}$$

where the modes have been split into those that are retained, $\boldsymbol{\Phi}_r$, and those that will be discarded, $\boldsymbol{\Phi}_d$. If $\mathbf{H}_w$ picks out those outputs that are assumed to be working (i.e. are elements of S_w), then we need to estimate $\mathbf{q}_{r,k} = \mathbf{q}_r(t_k)$ from

$$\mathbf{y}_{w,k} = \mathbf{H}_w\,\boldsymbol{\Phi}_r\,\mathbf{q}_{r,k} + \mathbf{H}_w\,\boldsymbol{\Phi}_d\,\mathbf{q}_{d,k} \tag{10}$$

where $\mathbf{y}_{w,k}$ denotes the response at the fully functioning sensors at time t_k and $\mathbf{q}_{d,k} = \mathbf{q}_d(t_k)$. Clearly the discarded modes in equation (10) could be neglected and the pseudo inverse used to estimate $\mathbf{q}_{r,k}$ from the resulting over-determined set of equations, as

$$\hat{\mathbf{q}}_{r,k} = \left(\mathbf{H}_w\,\boldsymbol{\Phi}_r\right)^{\dagger}\,\mathbf{y}_{w,k} \tag{11}$$

where $(\ )^{\dagger}$ denotes the usual Moore-Penrose pseudo inverse. This gives an estimate of the response at the functioning sensors as

$$\hat{\mathbf{y}}_{w,k} = \mathbf{H}_w\,\boldsymbol{\Phi}_r\left(\mathbf{H}_w\,\boldsymbol{\Phi}_r\right)^{\dagger}\,\mathbf{y}_{w,k} = \mathbf{P}\,\mathbf{y}_{w,k} \tag{12}$$

There will be an error introduced because

$$\left(\mathbf{H}_w\,\boldsymbol{\Phi}_r\right)^{\dagger}\,\mathbf{H}_w\,\boldsymbol{\Phi}_d \neq \mathbf{0} \tag{13}$$

and a better estimate may be obtained by using the orthogonality of the modes as

$$\hat{\mathbf{q}}_{r,k} = \left(\boldsymbol{\Phi}_r^T\mathbf{H}_w^T\mathbf{M}_{w,r}\mathbf{H}_w\boldsymbol{\Phi}_r\right)^{-1}\boldsymbol{\Phi}_r^T\mathbf{H}_w^T\mathbf{M}_{w,r}\,\mathbf{y}_{w,k} \tag{14}$$

where $\mathbf{M}_{w,r}$ is the mass matrix reduced to the degrees of freedom corresponding to the functioning sensors in the set S_w. Given that the mode shapes are assumed known, SEREP would be the most appropriate reduction method [15]. However, if the discarded modes lie outside the frequency range of interest then the estimator based on the pseudo inverse, equation (11), will be adequate. The corresponding estimate of the response is

$$\hat{\mathbf{y}}_{w,k} = \mathbf{H}_w\boldsymbol{\Phi}_r \left(\boldsymbol{\Phi}_r^T \mathbf{H}_w^T \mathbf{M}_{w,r} \mathbf{H}_w \boldsymbol{\Phi}_r\right)^{-1} \boldsymbol{\Phi}_r^T \mathbf{H}_w^T \mathbf{M}_{w,r}\, \mathbf{y}_{w,k} = \mathbf{P}\, \mathbf{y}_{w,k} \tag{15}$$

Both approaches give a projector matrix **P**, from the response space to the space of the lower modes. The quality of the model may be determined by reconstructing the response at the functioning sensors and producing the error as

$$\boldsymbol{\varepsilon}_{w,k} = (\mathbf{I} - \mathbf{P})\, \mathbf{y}_{w,k} \tag{16}$$

Reconstructing the responses of the faulty sensors gives the error as

$$\boldsymbol{\varepsilon}_{f,k} = \mathbf{y}_{f,k} - \mathbf{H}_f \boldsymbol{\Phi}_r \hat{\mathbf{q}}_{r,k} \tag{17}$$

where $\mathbf{H}_f$ picks out those outputs that are assumed to be faulty (i.e. are elements of S_f).

In practice we do not know which sensors are working and which are faulty. Therefore the errors in equations (16) and (17) are generated for all possible sets S_w and S_f. Of course the estimation of the modal participation factors has been performed at every time step, and so the errors will be produced at every time step. The average error over the time range of interest may be easily computed. The projector matrix, **P**, is constant for a particular choice of sets S_w and S_f and only needs to be computed once. Those sets where the error in the faulty sensor(s) is much greater than the error in the functioning sensors are then used to locate the faulty sensors.

5. VALIDATION VIA PRINCIPAL COMPONENT ANALYSIS

An alternative approach to estimating the modal participation factors is to estimate the modal model from the measured response. The mode shapes of this estimated model may then be compared to the analytical mode shapes using the co-ordinate modal assurance criterion (COMAC) [16], and the faulty sensors identified based on the degrees of freedom with the largest error. Some care must be exercised because data from the faulty sensors is being used to identify the modal model.

In fact it is not necessary to identify the experimental modes, since we are only interested in the subspace of the response, which is essentially a principal component analysis (PCA) [10-12]. Taking the singular value decomposition (SVD) of the response, for a block of data, gives,

$$\mathbf{Y} = [\mathbf{y}_{k+1}\, \mathbf{y}_{k+2} \cdots \mathbf{y}_{k+m}] = \mathbf{U}\, \mathbf{S}\, \mathbf{V}^T \tag{18}$$

where **U** and **V** are orthogonal and **S** is diagonal with the positive singular values arranged in decreasing order along the diagonal. If we were only interested in the subspace of the response at the functioning sensors then it would be computationally more efficient to calculate the SVD of

$$\mathbf{Y}\mathbf{Y}^T = \sum_{i=1}^{m} \mathbf{y}_{k+i}\, \mathbf{y}_{k+i}^T = \mathbf{U}\, \mathbf{S}^2\, \mathbf{U}^T \tag{19}$$

The assumption that only a limited number of modes influence the response may be checked by looking at the singular values. If we assume that the response contains r modes, then the response should be predominantly in the subspace defined by the first r columns of **U**, denoted $\mathbf{U}_r$. This subspace may then be compared to analytical subspace defined by the modal matrix $\boldsymbol{\Phi}_r$.

Thus far no allowance for the functioning or faulty sensors has been made. The subspaces defined by $\mathbf{H}_w\boldsymbol{\Phi}_r$ and $\mathbf{H}_w\mathbf{U}_r$ may be compared, requiring only a one SVD per data block. Alternatively the SVD of the response at the functioning

sensors, $\mathbf{H}_w\mathbf{Y}$, may be computed and the subspace obtained from this decomposition compared to $\mathbf{H}_w\mathbf{\Phi}_r$. Of course, the SVD of the response must then be computed for each sensor set, S_w. Thus, in this case, the response is projected onto the principal vectors, rather than the mode shapes. Suppose the subspace of the first r principal vectors of $\mathbf{H}_w\mathbf{Y}$ is $\mathbf{U}_{w,r}$. Then the resulting estimated response and projector matrix is (since $\mathbf{U}_{w,r}^T\mathbf{U}_{w,r} = \mathbf{I}_r$)

$$\hat{\mathbf{y}}_{w,k} = \mathbf{U}_{w,r}\,\mathbf{U}_{w,r}^T\,\mathbf{y}_{w,k} = \mathbf{P}\,\mathbf{y}_{w,k} \tag{20}$$

This only determines the model quality by comparing the consistency of the functioning sensors. The predicted response at the sensors in the faulty sensor set may be computed by estimating the transformation between $\mathbf{H}_w\mathbf{\Phi}_r$ and $\mathbf{U}_{w,r}$ and applying this same transformation to $\mathbf{H}_f\mathbf{\Phi}_r$, to give,

$$\hat{\mathbf{y}}_{f,k} = \mathbf{H}_f\mathbf{\Phi}_r\left(\mathbf{H}_w\mathbf{\Phi}_r\right)^{\dagger}\mathbf{U}_{w,r}\,\mathbf{U}_{w,r}^T\,\mathbf{y}_{w,k} \tag{21}$$

Thus, the residuals may be computed from the comparison of equations (20) and (21) with the measured response.

Notice that the mode shapes are used in the principal component analysis approach but only after the best response subspace of a given size has been computed. The subspaces may be conveniently compared using the concept of the angles between subspaces, which is a generalisation of the concept of angles between vectors. In three dimensions, it is easy to visualise the angle between a pair of lines, that is subspaces of dimension one, or indeed the angle between a line and a plane, that is subspaces of dimension one and two respectively. Bjorck and Golub [17] described the definition and calculation of the angles between subspaces. These ideas have been applied in structural dynamics in the areas of damage location [18], model updating using perturbed boundary condition testing [19], mode shape correlation [20] and sensor location [21]. The calculation gives as many angles as the smallest number of columns in the matrices defining the subspaces, and the largest angle gives a measure of how different the subspaces are. Unfortunately the subspace angles are quite insensitive to faults in individual sensors, and the residuals provide a more robust method to determine faulty sensors [22]. The approach adopted in the examples is to consider the statistical properties of the residuals.

6. STATISTICAL PROPERTIES OF RESIDUALS

The previous section outlined two methods to generate residuals for the two sets of sensors, $\boldsymbol{\varepsilon}_w$ and $\boldsymbol{\varepsilon}_f$. The data from the functioning sensors is essentially smoothed, by restricting the response to a subspace, determined either from the structure's lower mode shapes, or via a principal component analysis. Analysis of these residuals gives some idea of the noise properties of the measurements. The response at the faulty sensors is then predicted, and these residuals may be used to determine whether the hypothesis that this sensor set is faulty is likely to be true. The tool chosen to make these decisions is the statistical properties of the residuals.

In on line monitoring situations, the data will be taken in blocks of size m. The principal component analysis approach requires such a block, as shown in equations (18) and (19). The modal filtering may be performed at every sample independently, but this data will also be taken in blocks of size m. The mean and covariance of these data blocks may be estimated as [23]

$$\text{Mean}\left(\boldsymbol{\varepsilon}_w\right) = \bar{\boldsymbol{\varepsilon}}_w = \frac{1}{m}\sum_{i=1}^{m}\boldsymbol{\varepsilon}_{w,k+i} \tag{22}$$

and

$$\text{Var}\left(\boldsymbol{\varepsilon}_w\right) = \mathbf{V}_w = \frac{1}{m-1}\sum_{i=1}^{m}\left(\boldsymbol{\varepsilon}_{w,k+i} - \bar{\boldsymbol{\varepsilon}}_w\right)\left(\boldsymbol{\varepsilon}_{w,k+i} - \bar{\boldsymbol{\varepsilon}}_w\right)^T \tag{23}$$

and similarly for ε_f. Equipped with this information, providing assumptions are made concerning the probability density functions, then statistical inferences may be made from these quantities. Clearly the form of the measurement noise, the form of any modelling error, and the nature of any sensor fault will influence this distribution. Also, since the residuals have been obtained from all the measurements, it may well be that the residuals for individual sensors are correlated.

We wish to test whether the residuals have zero mean. The standard approach to hypothesis testing[23] is to calculate the test statistic

$$z_{w,i} = \frac{\bar{\varepsilon}_{w,i}}{\sqrt{\frac{\sigma_i^2}{m}}} \tag{24}$$

where $\bar{\varepsilon}_{w,i}$ is the i th element of $\bar{\varepsilon}_w$ and σ_i is the standard deviation of the i th sensor. If the probability distribution of the residuals is normal then $z_{w,i}$ will have a standard normal distribution with zero mean and unit variance. The m in equation (24) indicates that the standard deviation of the ensemble should reduce as more data points are added. Experience has shown in this application that this is not the case, and in the calculations that follow m is omitted in the calculation of the test statistic. If the standard deviation (or equivalently, the variance) is unknown then the estimate of the standard deviation from equation (23) is used $\left(\sigma_i^2 = \mathbf{V}_w(i,i)\right)$. $z_{w,i}$ will then have a t distribution with m-1 degrees of freedom, although for large m this will be approximately the same as the standard normal distribution. Because of the central limit theorem, the assumption that the probability density function will be normal is approximately satisfied. Thus the probability that the mean of the residual of the i th working sensors is zero is obtained from the standard normal distribution, $P_{w,i}$. Similarly, the probability that the mean of the residual of the i th faulty sensor is not zero, $P_{f,i}$, may be calculated. The probability that the split of sensors into faulty and working is the correct one is then

$$\mathrm{P}(\text{sensor split correct}) = \prod_{\text{sensor } i \in S_w} \mathrm{P}_{w,i} \prod_{\text{sensor } i \in S_f} \mathrm{P}_{f,i} \tag{25}$$

Although the actual probability should not be taken as an accurate measure, this process does provide a way to combine the data from all the residuals into a single figure of merit. These quantities may then be compared to establish which sensor groupings may be correct.

7. EXAMPLE

The proposed approaches will be tested on a simulated example of a cantilever beam excited at its tip, shown in Figure 1. The beam is 1 m long, 2.5 cm thick and 5 cm wide and the material properties of the beam are those of steel (a Young's modulus of 210 GN/m^2, a mass density of 7850 kg/m^3). The beam is simulated using 10 finite elements, and accelerometers are placed at every node position. The equations of motion are reduced by retaining only the lower 5 undamped modes in order to perform the numerical integration for the time response. Modal damping of 1% is included. Figure 2 shows the force applied at the tip, which simulates an impact excitation by a half sine pulse. Figure 3 shows the response at the beam tip, showing that the higher modes are excited by the force. Note the different time scales in Figures 2 and 3.

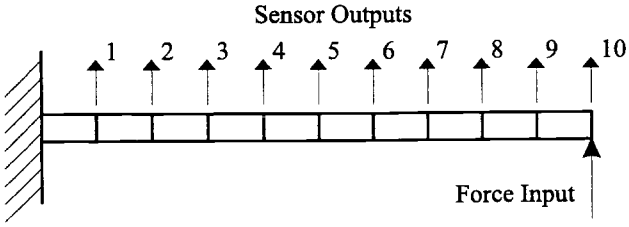

Figure 1. A schematic of the cantilever beam example

Two cases will be simulated. In both cases random noise is added to all the responses, and this noise is taken from a uniform distribution, which can be positive or negative, with a maximum magnitude of 2% of the maximum response from all the sensors. For case 1, a fault is introduced into sensor 6, such that the response has an additive error of 5% of the maximum response of all the sensors. Case 2 considers a multiplicative error, where the response of sensor 6 is multiplied by 1.5 (i.e. a 50% error). In all cases three modes of the structure (or three singular vectors) are used in the sensor validation checks.

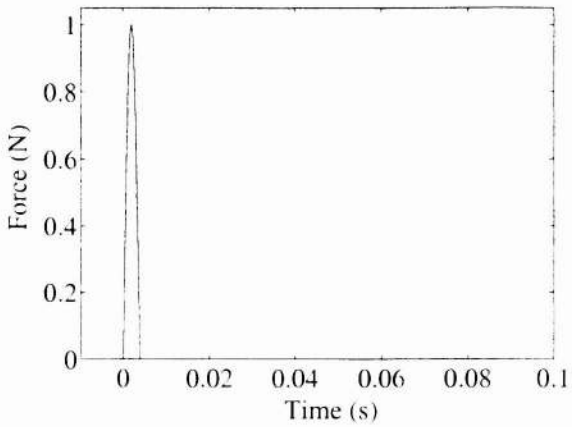

Figure 2. Half sine excitation at tip

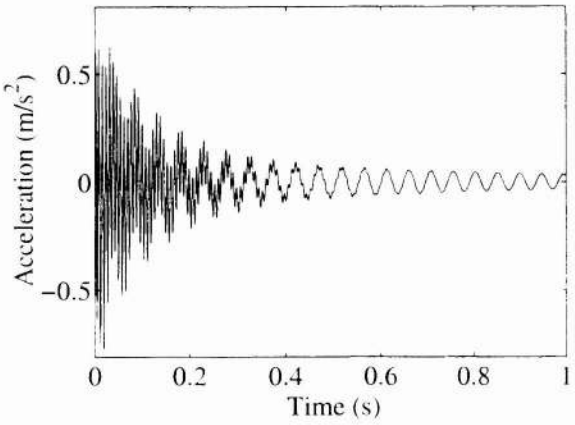

Figure 3. Response at cantilever beam tip

The modal filtering approach is first applied to case 1, using all 20000 time samples. Suppose that all sensors are assumed to be working. The residuals are computed and the mean and variance are shown in Figure 4, where the circles represent the mean for each sensor and the length of the error bars represents one standard deviation above and below the mean. It is clear that the residual corresponding to sensor 6 (the faulty sensor) is unlikely to have a zero mean. The faulty sensor has also caused non-zero mean residuals in nearby sensors. Indeed the probability given by equation (25) is 4×10^{-6}. Now suppose each sensor is assumed to be faulty in turn. Figure 5 shows the mean and standard deviation for each of these cases in turn using the same axis scaling. The sensor that is assumed to be faulty is given on each plot. Clearly when sensor 6 is assumed faulty all the working sensors produce residuals that are likely to have zero mean. Formally calculating the probabilities for each case gives the results shown in Figure 6, clearly indicating that sensor 6 is very likely to be faulty. If necessary groups of two sensors could be assumed faulty, and the probability that the sensor split into functioning and faulty sensor is correct could be calculated. However this is not done here.

The same calculations were performed on the data for case 2. Figures 7-9 show the results, and correspond to Figures 4 to 6 for case 1. Because the fault is a multiplicative error, it is the standard deviation that shows the fault most clearly. The probability that no sensor is faulty (corresponding to the mean and standard deviations in Figure 7) is 0.9951. Figure 9 shows that the probability that sensor 6 is faulty is the highest, but all these probabilities are very low. Thus, on the basis of this analysis, the conclusion would be that no sensor is faulty. It is clear that the multiplicative fault in case 2 is much more difficult to locate than the additive fault in case 1.

The principal component analysis approach was also tested on the data from cases 1 and 2. Figures 10 to 15 show the results and correspond to figures 4 to 9 for the modal filtering approach. Figure 10, where all sensors are assumed to be functioning, shows that the errors in the data due to the fault are spread amongst all the residuals. The probability that all sensors are functioning is 0.233, which is higher than the corresponding probability for the modal filtering approach. This is because the principal components are identified using the data from the faulty sensor and are able to allow for the faulty sensor in computing the best subspace of the response. This is also clearly shown in Figure 11, where the standard deviation of the sensors assumed to be functioning is very small and the standard deviation of the sensor assumed to be faulty is large. Computing the probabilities that each sensor in turn is faulty gives the results show in Figure 12, clearly indicating that sensor 6 is faulty. The PCA also has difficulty with the multiplicative fault of case 2, as shown in Figures 13 to 15, where the conclusion would again be that no sensors were faulty with a probability of 0.9997.

8. CONCLUSIONS

This paper has considered two approaches to determine if a sensor is faulty, both of which are based on considering the subspace of the response compared to the subspace generated by the lower modes of the structure. The approaches assume sets of sensors are faulty and the hypothesis that they are faulty (and the remainder are functioning) is tested. Clearly each potential subset of sensors should be tested. Using the residuals of the faulty sensor response, reconstructed from the functioning sensors and an associated model, seems to provide a method to locate faulty sensors. Residuals generated by

both methods outlined in this paper produce similar results on the simulated example, although if an accurate modal model is available the modal filtering approach performs better. Additive faults on the sensor were satisfactorily located, whereas multiplicative faults were much more difficult to detect. These multiplicative faults affected the variance of the residuals, and clearly hypothesis testing based on mean values is not suitable. It may be that alternative test statistics would be better for these types of faults, and this is the subject of further work.

ACKNOWLEDGEMENTS

Dr. Friswell gratefully acknowledges the support of the EPSRC through the award of an Advanced Fellowship. Prof. Inman gratefully acknowledges the support of National Science Foundation, grant no. CMS-9701471, through the Multidisciplinary Center for Earthquake Engineering Research, Buffalo, New York.

REFERENCES

1. D.W. Clark, "Sensor, Actuator, and Loop Validation," *IEEE Control Systems*, August 1995, pp. 39-45.
2. A.F. Dragoni, and P. Giorgini, "Sensor Validation for Nuclear Power Plants Through Bayesian Conditioning and Dempster's Rule of Combination," *Computers and Artificial Intelligence*, 17, pp. 151-168, 1998.
3. M. Henry, "Automatic Sensor Validation," *Control and Instrumentation*, 27, pp. 60-61, 1995.
4. K.E. Holbert, A.S. Heiger, and N.K. Alangrashid, "Redundant Sensor Validation by using Fuzzy-Logic," *Nuclear Science and Engineering*, 118, pp. 54-64, 1994.
5. S.C. Lee, "Sensor Value Validation based on Systematic Exploration of the Sensor Redundancy for Fault Diagnosis KBS," *IEEE Transactions on Systems Man and Cybernetics*, 24, pp. 594-605, 1994.
6. M.R. Napolitano, D.A. Windon, J.L. Casanova, M. Innocenti, and G. Silvestri, "Kalman Filters and Neural-Network Schemes for Sensor Validation in Flight Control Systems," *IEEE Transactions on Control Systems Technology*, 6, pp. 596-611, 1998.
7. E.A. Scarl, J.R. Jamieson, and C.I. Delaune, "Diagnosis and Sensor Validation Through Knowledge of Structure and Function," *IEEE Transactions on Systems Man and Cybernetics*, 17, pp. 360-368, 1987.
8. J.C.-Y. Yang, and D.W. Clarke, "Control using Self-Validating Sensors," *Transactions of the Institute of Measurement and Control*, 18, pp. 15-23, 1996.
9. N.D. Walker, and G.F. Wyattmair, "Sensor Signal Validation using Analytical Redundancy for a Cold-Rolling Mill," *Control Engineering Practice*, 3, pp. 753-760, 1995.
10. R. Dunia, S.J. Qin, T.F. Edgar, and T.J. McAvoy, "Identification of Faulty Sensors using Principal Component Analysis," *AIChE Journal*, 42, pp. 2797-2812, 1996.
11. R. Dunia, and S.J. Qin, "Subspace Approach to Multidimensional Fault Identication and Reconstruction," *AIChE Journal*, 44, pp. 1813-1831, 1998.
12. S.J. Qin, H. Yue, and R. Dunia, "Self-Validating Inferential Sensors with Application to Air Emission Monitoring," *Industrial & Engineering Chemistry Research*, 36, pp. 1675-1685, 1997.
13. A.J. Millar, *Subset Selection in Regression*, Chapman and Hall, 1990.
14. M.I. Friswell, J.E. Mottershead, and H. Ahmadian, "Combining Subset Selection and Parameter Constraints in Model Updating," *Journal of Vibration and Acoustics*, 120, pp. 854-859, 1998.
15. J.C. O'Callahan, P. Avitabile, and R. Riemer, "System Equivalent Reduction Expansion Process (SEREP)," *7th IMAC*, Las Vegas, pp. 29-37, February 1989.
16. N.A.J. Lieven, and D.J. Ewins, "Spatial Correlation of Mode Shapes, the Coordinate Modal Assurance Criterion," *6th IMAC*, Kissimmee, Florida, pp. 690-695, 1988.
17. A. Bjorck, and G.H. Golub, "Numerical Methods for Computing Angles Between Linear Subspaces," *Mathematics of Computation*, 27, pp. 579-594, 1973.
18. A.P. Cherng, and M.K. Abelhamid, "A Signal Subspace Correlation (SSC) Index for Detection of Structural Changes," *11th IMAC*, Orlando, Florida, pp. 232-239, 1993.
19. M. Yang, and D. Brown, Model Updating Techniques using Perturbed Boundary Condition (PBC) Testing Data, *14th IMAC*, Detroit, Michigan, pp. 776-782, 1996.
20. S.D. Garvey, J.E.T. Penny, and M.I. Friswell, "Quantifying the Correlation Between Measured and Computed Mode-Shapes," *Journal of Vibration and Control*, 2, pp. 123-144, 1996.

21. S.D. Garvey, M.I. Friswell, and J.E.T. Penny, "Evaluation of a Strategy for Optimal Choice of Measurement Locations Based on the Concept of Minimum Angles Between Subspaces," *14th IMAC*, Dearborn, Michigan, pp. 1546-1552, 1996.
22. M.I. Friswell, and D.J. Inman, "Sensor Validation for Smart Structures. *18th IMAC,* San Antonio, Texas, pp. 483-489, 2000.
23. R.J. Freund, and W.J. Wilson, *Regression Analysis: Statistical Modelling of a Response Variable*, Academic Press, 1998.

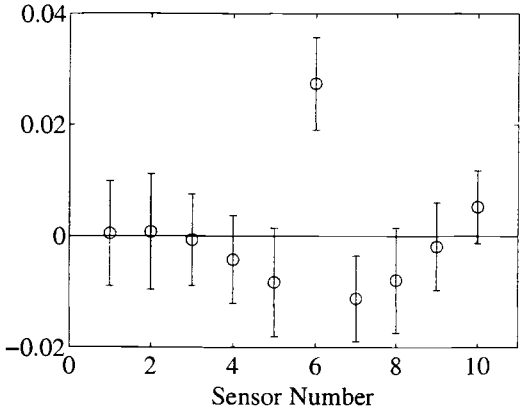

Figure 4. The mean (circles) and standard deviation (error bars) of the residuals for each sensor for fault case 1 using the modal filtering approach. All sensors assumed functioning

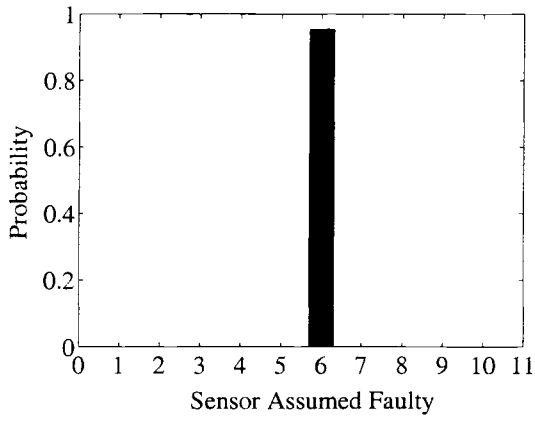

Figure 6. Probabilities for one sensor in turn assumed faulty for fault case 1 using the modal filtering approach

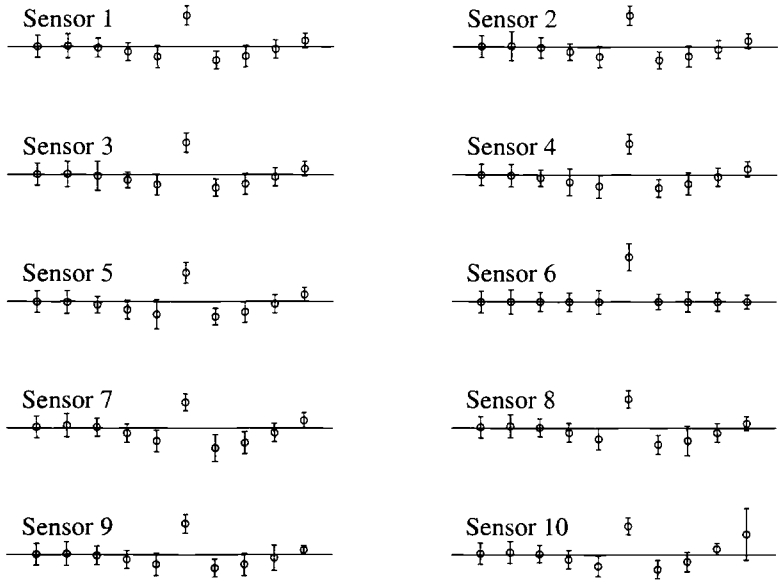

Figure 5. The mean (circles) and standard deviation (error bars) of the residuals for each sensor for fault case 1 using the modal filtering approach. One named sensor assumed faulty in each plot

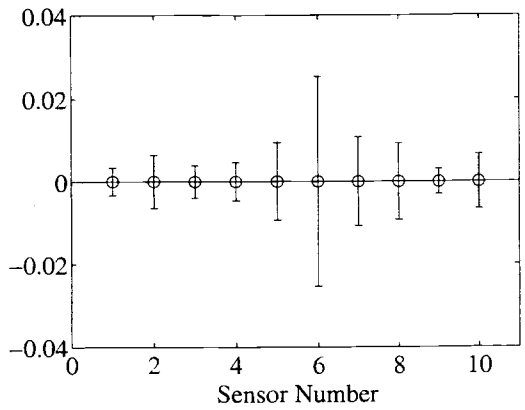

Figure 7. The mean (circles) and standard deviation (error bars) of the residuals for each sensor for fault case 2 using the modal filtering approach. All sensors assumed functioning

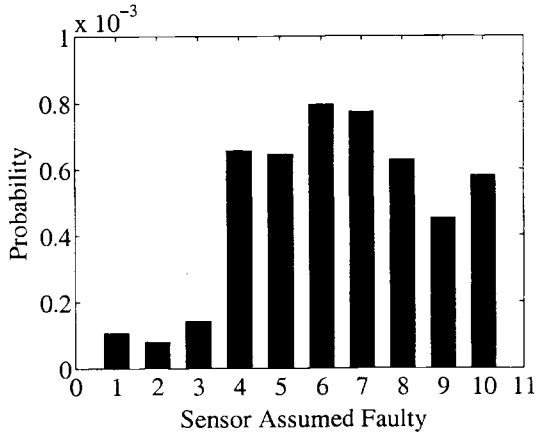

Figure 9. Probabilities for one sensor in turn assumed faulty for fault case 2 using the modal filtering approach

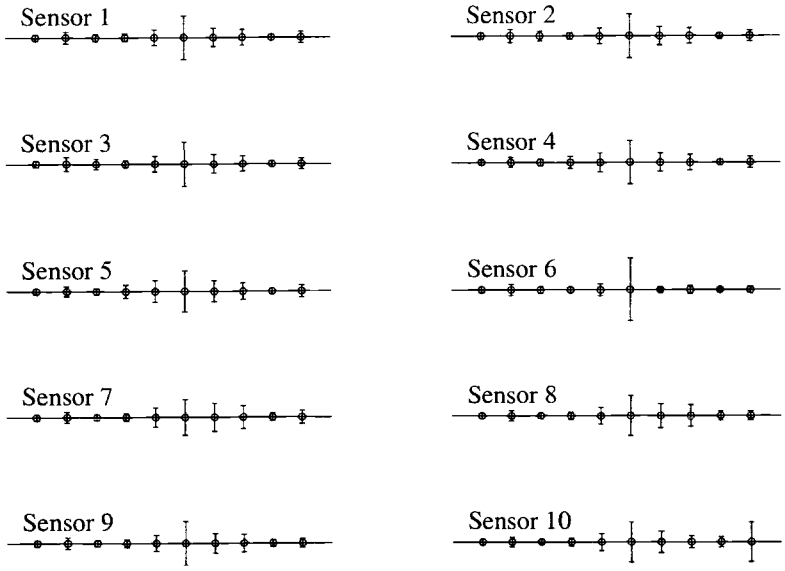

Figure 8. The mean (circles) and standard deviation (error bars) of the residuals for each sensor for fault case 2 using the modal filtering approach. One named sensor assumed faulty in each plot

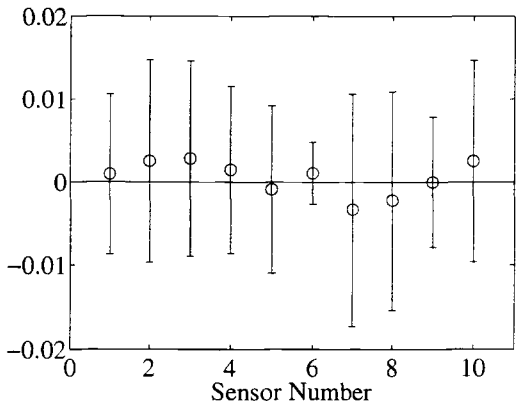

Figure 10. The mean (circles) and standard deviation (error bars) of the residuals for each sensor for fault case 1 using the PCA approach. All sensors assumed functioning

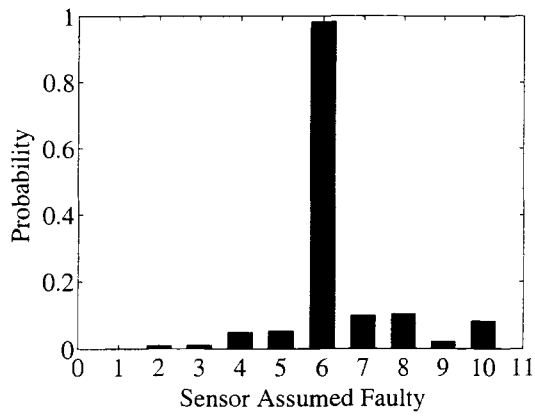

Figure 12. Probabilities for one sensor in turn assumed faulty for fault case 1 using the PCA approach

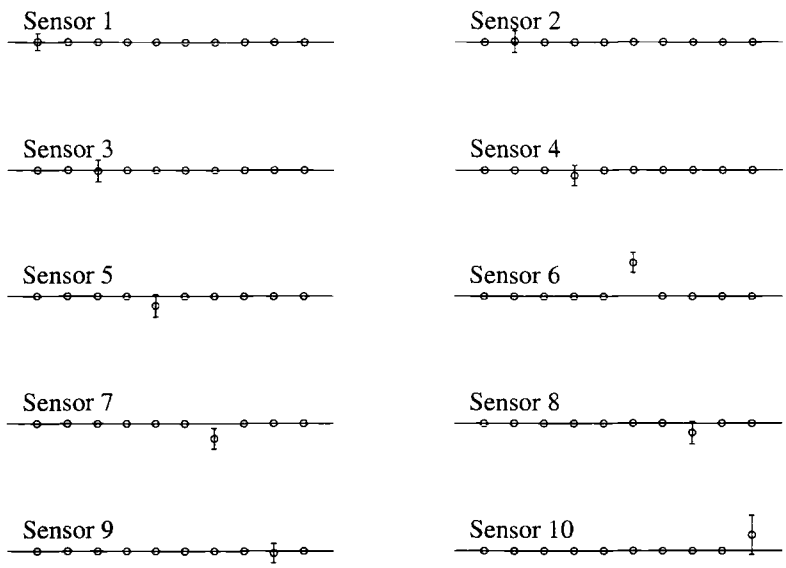

Figure 11. The mean (circles) and standard deviation (error bars) of the residuals for each sensor for fault case 1 using the PCA approach. One named sensor assumed faulty in each plot

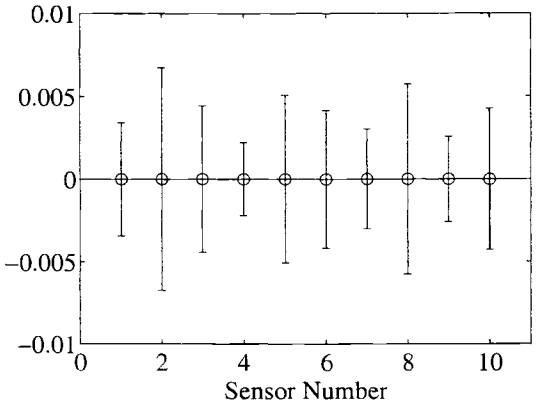

Figure 13. The mean (circles) and standard deviation (error bars) of the residuals for each sensor for fault case 2 using the PCA approach. All sensors assumed functioning

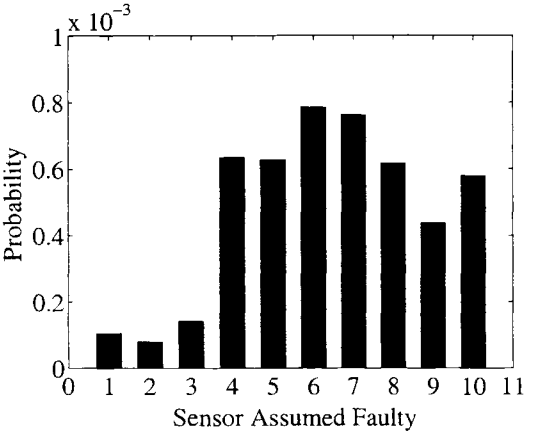

Figure 15. Probabilities for one sensor in turn assumed faulty for fault case 2 using the PCA approach

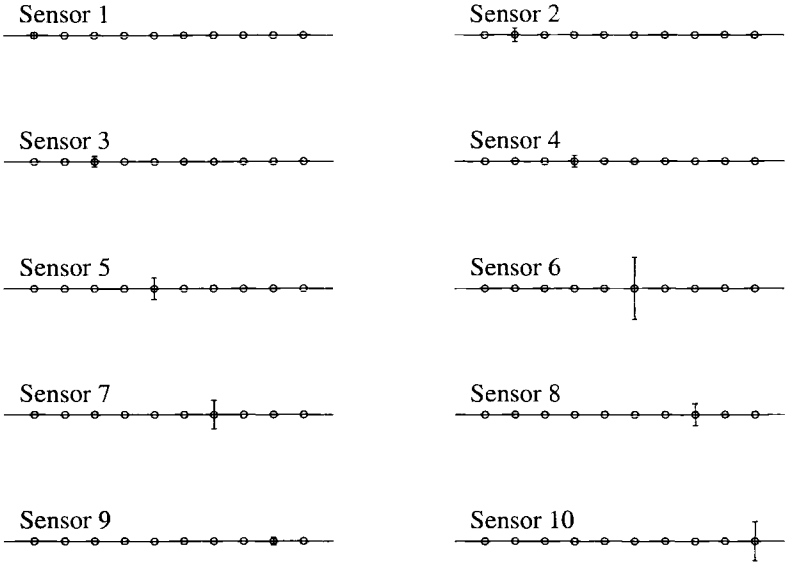

Figure 14. The mean (circles) and standard deviation (error bars) of the residuals for each sensor for fault case 2 using the PCA approach. One named sensor assumed faulty in each plot