

**OPTIMALLY ROBUST REDUNDANCY RELATIONS
FOR FAILURE DETECTION IN UNCERTAIN SYSTEMS⁺**

Xi-Cheng Lou*
Alan S. Willsky*
George C. Verghese**

Room 35-233
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 01239

Abstract

All failure detection methods are based, either explicitly or implicitly, on the use of redundancy, i.e. on (possibly dynamic) relations among the measured variables. The robustness of the failure detection process consequently depends to a great degree on the reliability of the redundancy relations, which in turn is affected by the inevitable presence of model uncertainties. In this paper we address the problem of determining redundancy relations that are optimally robust, in a sense that includes several major issues of importance in practical failure detection, and that provides a significant amount of intuition concerning the geometry of robust failure detection. We also give a procedure, involving the construction of a single matrix and its singular value decomposition, for the determination of a complete sequence of redundancy relations, ordered in terms of their level of robustness. This procedure also provides the basis for comparing levels of robustness in redundancy provided by different sets of sensors.

⁺This research was supported in part by the Office of Naval Research under Grant N00014-77-C-0224, by NASA Ames and NASA Langley Research Centers under Grant NGL-22-009-124, and by the Air Force Office of Scientific Research under Grant AFOSR-82-0258.

*Also affiliated with the M.I.T. Laboratory for Information and Decision Systems.

**Also affiliated with the M.I.T. Laboratory for Electromagnetic and Electronic Systems.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE MAR 1985		2. REPORT TYPE		3. DATES COVERED 00-03-1985 to 00-03-1985	
4. TITLE AND SUBTITLE Optimally Robust Redundancy Relations for Failure Detection in Uncertain Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology, Laboratory for Information and Decision Systems, 77 Massachusetts Avenue, Cambridge, MA, 02139				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 33	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1. Introduction

A wide variety of techniques has been proposed in recent years for the detection, isolation, and accommodation of failures in dynamic systems (see, for example, the surveys in [1,4]). In one way or another, all of these methods involve the generation of signals that are accentuated by the presence of particular failures if these failures have actually occurred. The procedures for generating these signals in turn depend on models relating the measured variables. Consequently, if any errors in these models have effects on the observables that are at all like the effects of any of the failure modes, then these model errors may also accentuate the signals. This leads us directly to the issue of robust failure detection, that is, the design of a system that is maximally sensitive to the effects of failures and minimally sensitive to model errors.

The work described here focuses on directly designing a failure detection system that is insensitive to model errors (rather than designing a system that attempts to compensate the detection algorithm by estimating uncertainties on-line, see [6, 7, 12]). The initial impetus for our approach came from the work reported in [5, 13], in the context of aircraft failure detection. The noteworthy feature of that project was that the dynamics of the aircraft were decomposed in order to analyze the relative reliability of each individual source of potentially useful failure detection information. In this way, a design was developed that utilized only the most reliable information.

In [2] we presented the results of our initial attempt to extract the essence of the method used in [9, 13] in order to develop a general approach to robust failure detection. As discussed in those references and in others (such as [3, 7, 8]), all failure detection systems are based on exploiting analytical redundancy relations or (generalized) parity checks. These are simply functions of the temporal histories of the measured quantities that have the property of being small (ideally zero) when the system is operating normally. Essentially all of the recently developed general approaches to failure detection make implicit, rather than explicit use of all of these relations. That is, these general methods use an overall dynamic model as the basis for designing failure detection algorithms. While such a model certainly captures all of the relationships among the measured variables, it does not in any way discriminate among these individual relationships. For this reason, a top-down application of any of these methods mixes together information of varying levels of reliability. What would clearly be preferable would be a general method for explicitly

identifying and utilizing only the most reliable of the redundancy relations.

One criterion for measuring the reliability of a particular redundancy relation was presented in [2] and was used to pose an optimization problem to determine the most reliable relation. This criterion has the feature that it specifies robustness with respect to a particular operating point, thereby allowing the possibility of adaptively choosing the best relations. However, a drawback of this approach is that it leads to an extremely complex optimization problem. Moreover, if one is interested in obtaining a list of redundancy relations that is ordered from most to least reliable, one must essentially solve a separate optimization problem for each relation in the list.

In this paper we look at an alternative measure of reliability for a redundancy relation. Not only does this alternative have a helpful geometric interpretation, but it also leads to a far simpler optimization procedure, involving a single singular value decomposition. In addition, it allows us in a natural and computationally feasible way to consider issues such as scaling, relative merits of alternative sensor sets, and explicit tradeoffs between detectability and robustness.

In Section 2 we review the notion of analytical redundancy for perfectly known models, and then provide a geometric interpretation that forms the starting point for our investigation of robust failure detection. Section 3 addresses the problem of robustness using our geometric ideas, and solves a version of the optimally robust redundancy problem. In Section 4 we discuss extensions to include three important issues not included in Section 3: noise, known inputs, and the detection/robustness tradeoff. We conclude the paper in Section 5 with a discussion of several other topics, including the relationship of our results to those in [2] and the use of this formalism to measure and compare the levels of robust redundancy associated with different system configurations.

2. Redundancy Relations

This paper focuses attention on linear, time-invariant, discrete-time systems. In this section we consider the uncertainty-free model

$$x(k+1) = Ax(k) + Bu(k), \quad (1)$$

$$y(k) = Cx(k) + Du(k), \quad (2)$$

where x is an n -dimensional state vector, u is an m -dimensional vector of known inputs, y is an r -dimensional vector of measured outputs, and A , B , C and D are known matrices of appropriate dimensions. A redundancy relation for this model is some linear combination of present and lagged values of u and y that is identically zero if no changes (i.e. failures) occur in (1), (2).

As discussed in [2], redundancy relations can be specified mathematically in the following way. The subspace of $(s+1)r$ -dimensional vectors given by

$$P = \left\{ v \mid v^T \begin{bmatrix} C \\ CA \\ \vdots \\ CA^s \end{bmatrix} = 0 \right\} \quad (3)$$

is called the parity space of order s (to be distinguished from the s -step unobservable subspace, which corresponds to the right null space of the matrix in (3) rather than its left null space). We shall denote $(s+1)r$ by N . Every vector v in (3) can be associated at any time k with a parity check, $r(k)$:

$$r(k) = v^T \left[\begin{bmatrix} y(k-s) \\ y(k-s+1) \\ \vdots \\ y(k) \end{bmatrix} - H \begin{bmatrix} u(k-s) \\ u(k-s+1) \\ \vdots \\ u(k) \end{bmatrix} \right], \quad (4)$$

$$H = \begin{bmatrix} D & & & & \\ CB & D & & & 0 \\ CAB & CB & D & & \\ CA^2B & CAB & CB & D & \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ CA^{s-1}B & \vdots & \vdots & CAB & CB & D \end{bmatrix}. \quad (5)$$

(The development in Sections 2 to 4 deals with a single, fixed value of s . Therefore, to avoid notational clutter, we shall not index subspaces such as P in (3) or matrices such as H in (4) with the subscript s . Consideration of different values of s is contained in Section 5.) By (1), (2), the quantity in brackets [.] in (4) equals

$$\begin{bmatrix} C \\ CA \\ \vdots \\ CA^s \end{bmatrix} x(k-s) \quad . \quad (6)$$

Hence, by (3), we see that the simple redundancy relation or parity check

$$r(k) = 0 \quad (7)$$

is satisfied.

It is evident from (4) and (7) that a redundancy relation is simply an input-output model for (or constraint on) part of the dynamics of the system (1), (2). This interpretation of a redundancy relation allows us to make contact with the numerous existing failure detection methods. These methods are typically based on a noisy version of the model (1), (2) that represents normal system behavior, together with a set of deviations from this model that represent the several failure modes. However, rather than applying such methods to a single, all-encompassing model as in (1), (2), one could alternatively apply the same techniques to individual models as in (4), (7), or to a combination of several of these, which serves to isolate individual (or specific groups of) parity checks. (See Section 5 for some further comments on this point.) This is precisely what was done in [5, 13], for example. The advantage of such an approach is that it allows one to separate the information provided by redundancy relations of differing levels of reliability, something that is not easily done when one starts with the overall model (1), (2), which combines all redundancy relations.

In the next two sections we address the main problem of this paper, which is the determination of optimally robust redundancy relations. The key to this approach is obtained by re-examining (3)-(7), in order to suggest a geometrical interpretation of parity relations. In particular, consider the model (1), (2) and let Z denote the range of the matrix in (3). Then the parity space P is the orthogonal complement of Z , and a complete set of parity checks, of order s and of the form (4), (7), is given by the orthogonal projection of the vector of input-

adjusted observations

$$\begin{bmatrix} y(k-s) \\ y(k-s+1) \\ \vdots \\ y(k) \end{bmatrix} - H \begin{bmatrix} u(k-s) \\ u(k-s+1) \\ \vdots \\ u(k) \end{bmatrix} \quad (8)$$

onto P .

To illustrate this, consider an example in which the first two components of y measure scaled versions of the same variable, i.e.

$$y_2(k) = ay_1(k) . \quad (9)$$

Then, as illustrated in Figure 1, the subspace Z in $y_1 - y_2$ space is simply the line specified by (9). Furthermore, in this case the obvious parity relation is

$$r(k) = y_2(k) - ay_1(k) , \quad (10)$$

which is nothing more than the orthogonal projection of the observed pair of values $y_1(k)$ and $y_2(k)$ onto the line P perpendicular to Z (Figure 1). For interpretations of the space P in purely matrix terms and in terms of polynomial matrices, we refer the reader to [9] and [3], respectively. It is the geometric interpretation, however, that we shall utilize here.




Figure 1: An Example of the Geometric Interpretation of Parity Relations.

3. A Geometric Approach to Robust Redundancy

To begin, let us focus on a model that is not driven by either unknown noise or known signals:

$$x(k+1) = A_q x(k) \quad (11)$$

$$y(k) = C_q x(k) \quad (12)$$

where q indexes the models associated with different possible values of the unknown parameters. Throughout this paper (except for a brief discussion in Section 5), we consider only the case where q is taken from a finite set of possibilities, say $q=1, 2, \dots, Q$. In practice, this might involve choosing representative points out of the actual, continuous range of parameter values, reflecting any desired weighting on the likelihood or importance of particular sets of parameter values.

Define the (s -step) observation space Z_q by

$$Z_q = \text{range} \begin{bmatrix} C_q \\ C_q A_q \\ \vdots \\ C_q A_q^s \end{bmatrix}. \quad (13)$$

This is the subspace in which the window of observations for the system (11), (12) lives, as $x(k-s)$ varies over all possible values. For a given q , the parity space is the orthogonal complement, P_q , of Z_q . However, the orthogonal complement of one observation space will not be the orthogonal complement of another distinct observation space. It is therefore in general impossible to find parity checks that are perfect for all possible values of q . That is, in general we cannot find a subspace P that is orthogonal to Z_q for all q .

What would seem to make sense in this case is to choose a subspace P that is "as orthogonal as possible" to all possible Z_q . Returning to our simple example, suppose that $y_2 = ay_1$ but that ' a ' is only known to lie in some interval. In this case we obtain the picture shown in Figure 2. The shaded regions here represents the range of (y_1, y_2) values consistent with the uncertainty in ' a '. Intuitively, what would seem to be a good choice for P (assuming that ' a ' is equally likely to lie anywhere in the interval (24)) is the line that bisects the obtuse angle between the shaded sectors in Figure 2. It is precisely this geometric picture that is generalized and built upon in this paper.

For the general case, our procedure will be to first compute an average observation space Z_0 that is as close as possible, in a sense to be made precise, to all of the Z_q . We shall then choose P to be the orthogonal complement of Z_0 . (This idea is also illustrated in Figure 2, where the average observation space Z_0 is depicted as the line that bisects the shaded region, and the line P then represents its orthogonal complement.) Note that the Z_q are subspaces of possibly differing dimensions, embedded in a space of dimension $N = (s+1)r$, corresponding to histories of the last $s+1$ values of the r -dimensional output. Consequently, if we would like to determine the p best parity checks (so that $\dim P = p$), we need to find a subspace Z_0 of dimension $N-p$.

A Preliminary Scaling: Before stating the criterion that defines Z_0 , it is necessary to take account of a fact that has been glossed over so far. It is not sufficient to simply examine the subspaces in which signals lie; one has also to consider the characteristic magnitudes and directions of the excursions of signals in the subspaces to which they are confined. It will typically be the case that some components (or combinations of components) of $x(k-s)$ are larger than others, because they may be measured in different units and excited differently. Hence certain excursions in observation space are more likely than others. To take account of this, assume for now that we are able to find a nonsingular scaling matrix M_q such that, with the change of basis

$$x = M_q w, \quad (14)$$

one obtains a variable w that is governed by a similarity-transformed version of (11), (12) and has "equally likely" excursions of "unit length" in each direction under the q -th model. This sort of normalization is discussed more at the end of this section and in Section 4.1, where observation and process noise are incorporated into the model. (See also [11], in which scaling is also considered in the context of the design of a failure detection system.) We can now use the columns of the matrix

$$\begin{bmatrix} C_q \\ C_q A_q \\ \vdots \\ C_q A_q^s \end{bmatrix} M_q \quad (15)$$

as a spanning set for Z_q . We shall denote the matrix in (15) by the non-boldface Z_q . We shall, in the remainder of this paper, consistently use a **boldface** capital letter to denote the subspace spanned by the columns of a matrix that is denoted by the corresponding non-boldface capital.




Figure 2: Illustrating the choice of G in the presence of uncertain parameters.

The criterion for the best choice of Z_0 may now be defined in the following manner. With Z_1, \dots, Z_Q denoting the scaled matrices in (15) whose columns span the possible subspaces in which the observation histories may lie under normal conditions, define the $N \times Qn$ matrix

$$Z = [Z_1: \dots : Z_Q] \quad (16)$$

The optimum choice for Z_0 is then taken to be the span of the columns of the matrix Z_0 that minimizes

$$\|Z - Z_0\|_F^2, \quad (17)$$

subject to the constraint that $\text{rank } Z_0 = N-p$ (which ensures that the orthogonal complement P of Z_0 has dimension p). Here $\|\cdot\|_F$ denotes the Frobenius norm, which is defined as the sum of the squares of the entries of the associated matrix. The matrix Z_0 is thus chosen so that the sum of the squared distances between the columns of Z and of Z_0 is minimized, subject to the constraint that Z_0 contains only $N-p$ linearly independent columns.

The optimization problem we have just posed is easy to solve. In particular let the singular value decomposition (see [14, 15]) of Z be given by

$$Z = U \Sigma V \quad (18)$$

where

$$\Sigma = \begin{bmatrix} \sigma_1 & & & & & \\ & \sigma_2 & & & & \\ & & \ddots & & & \\ & & & 0 & & \\ 0 & & & & \ddots & \\ & & & & & \sigma_N \\ & & & & & & 0 \end{bmatrix}, \quad (19)$$

and U and V are orthogonal matrices. Here $\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_N$ are the singular values of Z , ordered by magnitude. Note that we have actually assumed $N \leq Qn$. If this is not the case, we can make it so without changing the optimum choice of Z_0 by padding Z with additional columns of zeros. As shown in [17] (see also [18]), the matrix Z_0 minimizing (17) is given by

$$Z_0 = U \left[\begin{array}{ccc|c} 0 & & & \\ & 0 & & \\ & & 0 & \\ & & \sigma_{p+1} & \\ 0 & & \cdot & \\ & & & \sigma_N \\ & & & \end{array} \right] V . \quad (20)$$

Moreover, since the columns of U are orthonormal, we immediately see that the orthogonal complement of the range Z_0 of Z_0 is given by the first p left singular vectors of Z_0 , i.e. the first p columns of U . Consequently, an orthonormal basis for the parity space P is given by

$$P = [u_1, \dots, u_p] \quad (21)$$

and u_1, \dots, u_p define optimum redundancy relations or parity checks.⁺

There are additional reasons for choosing this method for determining Z_0 and P , apart from the fact that the computation just described is quite straightforward. Firstly, minimization of the criterion in (17) does produce a space that is as close as possible in a natural sense to a specified set of directions, namely the columns of $\{Z_q, q = 1, \dots, Q\}$. Thanks to the scaling (14), these columns represent a complete set of "equally likely" directions in the observation space Z_q (corresponding to the "equally likely" values of the scaled state $w = [1, 0, \dots, 0]^T$, $[0, 1, \dots, 0]^T$, etc.). A second (and more precisely stated) reason follows from an alternative interpretation of our choice of P that provides some very useful insight.

Specifically, recall that what we wish to do is to find a subspace P that is as orthogonal as possible to all the subspaces Z_q . Translating this to statements about bases for these spaces, we would like to choose an $N \times p$ matrix P , normalized by the condition that it have orthonormal columns (i.e. $P^T P = I_p$, so that P is the orthogonal projection onto the subspace P), to make each of the matrices $P^T Z_q$ as close to zero as possible. Now, as shown in the Appendix, the choice of P given in (21) also minimizes

$$J = \sum_{q=1}^Q \|P^T Z_q\|_F^2, \quad (22)$$

yielding the minimum value

⁺Note that if $\sigma_{p+1} = 0$, then (a) Z_0 actually has rank less than $N-p$ and (b) there is a perfectly robust parity space of dimension at least $p+1$.

$$J^* = \sum_{i=1}^P \sigma_i^2 . \quad (23)$$

In fact, as illustrated in the Appendix, the same choice of P can also be shown to minimize other physically meaningful criteria.

Some important points about the result (22), (23) should be noted. To begin with, one can now see a straightforward way in which to include unequal weightings on each of the terms in (22). Specifically, if a_q are positive numbers, then minimizing

$$J_1 = \sum_{q=1}^Q a_q \| P^T Z_q \|^2_F \quad (24)$$

is accomplished using the same procedure described previously, but with Z_q replaced by $\sqrt{a_q} Z_q$. Carrying this one step further, if we normalize the a_q so that they sum to one, we can think of them as representing the prior probabilities for each of the possible system models. Thus J_1 in (24) can be interpreted as the expected value of $\| P^T Z_q \|^2_F$, where the expectation is taken over the model uncertainty. Furthermore, if we interpret the scaling (14) as producing a state w with unit covariance (i.e. $E[ww^T] = I$), then $\| P^T Z_q \|^2_F$ can be interpreted as $E_q(\| r(k) \|^2)$, where $r(k)$ now (unlike in (4)) is being used to denote the vector whose entries are the complete set of parity checks determined by the projection P ,

$$r(k) = P^T \begin{bmatrix} y(k-s) \\ y(k-s+1) \\ \vdots \\ y(k) \end{bmatrix} = P^T Z_q w(k-s) , \quad (25)$$

and E_q represents the expectation over $w(k-s)$, assuming that the data is generated by the q -th model. Combining this with the probabilistic interpretation of the a_q , we see that

$$J_1 = E(\| r(k) \|^2) , \quad (26)$$

where E denotes expectation over $w(k-s)$ and the model uncertainty. It is on this interpretation that we build in the next section.

Finally, note that the optimum value (23) provides us with an

interpretation of the singular values as measures of robustness and provides a sequence of parity relations ordered from most to least robust: u_1 is the most reliable parity relation, with σ_1^2 as its robustness measure; u_2 is the next best relation, with σ_2^2 as its robustness measure; etc. Consequently, from a single singular value decomposition, we can obtain a complete solution to the robust redundancy relation problem for a fixed value of s , i.e. for a fixed-length time history of output values.

4. Three Extensions

In this section we develop three extensions of the result of the preceding section, through modifications that entail no fundamental increase in complexity. The treatment of noise is first addressed, in Section 4.1, while the inclusion of known inputs is discussed in Section 4.2. Finally, the issue of designing parity checks for robust detection of a particular failure mode is examined in Section 4.3.

4.1 Observation and Process Noise

In addition to choosing parity relations that are maximally insensitive to model uncertainties, it is also important to choose relations that suppress noise. Consider the model

$$x(k+1) = A_q x(k) + B_q u(k), \quad (27)$$

$$y(k) = C_q x(k) + D_q u(k), \quad (28)$$

where $u(\cdot)$ is a zero mean, unit covariance, white noise process. We assume that x and y have attained stationarity, and that the steady-state covariance of x is given by

$$S_q = M_q M_q^T \quad (29)$$

The time window of observations for (27), (28) is now given by

$$\begin{bmatrix} y(k-s) \\ y(k-s+1) \\ \vdots \\ y(k) \end{bmatrix} = \begin{bmatrix} C_q \\ C_q A_q \\ \vdots \\ C_q A_q^s \end{bmatrix} M_q w(k-s) + H_q \begin{bmatrix} u(k-s) \\ u(k-s+1) \\ \vdots \\ u(k) \end{bmatrix} \quad (30)$$

where $w(k-s)$ has zero mean and unit covariance -- cf. (14), (15) and the discussion at the end of Section 3 -- and H_q has the same structure as in (8), except that all matrices are replaced by their subscripted versions, since it is the q -th model that is under consideration. We shall write (30) more compactly as

$$Y(k) = Z_q w(k-s) + H_q U(k), \quad (31)$$

with the definitions of the symbols being obvious from (30). In particular, note that the $U(k)$ has unit covariance and is independent of $w(k-s)$.

A natural extension of the minimization criterion (24), (26) is then provided by

$$J = \sum_{q=1}^Q a_q E_q(\|r(k)\|^2) \quad (32)$$

where

$$r(k) = P^T Y(k) \quad (33)$$

and where E_q denotes the expectation over $w(k-s)$ and $U(k)$, assuming that the data is generated by the q -th model. As before, J is to be minimized by choice of P that satisfies $P^T P = I$, and the parity space P will then be taken to be the range of P .

For simplicity, let us first assume that $a_q = 1$ for all q . It is then quite directly seen that

$$\begin{aligned} J &= \sum_{q=1}^Q \text{tr}[P^T (Z_q Z_q^T + H_q H_q^T) P] \\ &= \sum_{q=1}^Q \|P^T [Z_q : H_q]\|_F^2. \end{aligned} \quad (34)$$

From this it is evident, given our previous results, that the optimum choice of P is computed by performing a singular value decomposition on the matrix

$$T = [Z_1 : H_1 : \dots : Z_Q : H_Q] \quad (35)$$

If the a_q are not all identical, then we simply modify T by scaling Z_q and H_q by $\sqrt{a_q}$.

It is evident from the above that the effect of noise is simply to define additional directions to which the columns of P should be as orthogonal as possible. That is, P is to be chosen so that the parity check $r(k)$ has minimal response both to the likely sequences of values of the ideal noise-free observations (as specified by the columns of Z_q) and to the directions in which the observation noise and process noise have their maximum effects (as determined by the columns of H_q). The solution of this problem yields, as before, a complete set of parity

checks, corresponding to the left singular vectors of T , ordered in terms of their degrees of insensitivity to model errors and noise (as measured by the corresponding singular values).

4.2 Known Inputs

The analysis of the preceding section can be modified somewhat to allow us to consider the case in which some of the driving terms in (27) are known inputs. To simplify the discussion in this section, we assume that all of the components of $u(k)$ are known inputs. The extension to the case when there are both known inputs and noise is straightforward.

The key difference between the case in which $u(k)$ is unmeasured and the case in which it is measured is that in the latter case we can adjust the measured outputs $y(k)$ to account for the effect of the measured inputs $u(k)$ (see the discussion in Section 2). That is, we can consider defining a vector of parity checks of the form

$$r(k) = P^T \begin{bmatrix} Y(k) \\ U(k) \end{bmatrix} \quad (36)$$

where $P^T P = I_p$. The question then is, how do we measure the robustness of $r(k)$. Clearly, since $U(k)$ is known, we can consider defining a robustness measure relative to any specified input sequence $U(k)$. This approach is closer to the spirit of the work of Chow and Willsky [2]. As discussed in Section 5, such an approach allows one to adjust the parity matrix P on-line by (in effect) scheduling it with respect to $U(k)$, but the price that is paid for this is significantly greater on-line and off-line computational complexity.

What we shall do instead is to follow the same philosophy we have used upto this point. That is, we shall attempt to find a single matrix P that minimizes the norm of $r(k)$ on the average, as $w(k-s)$ and $U(k)$ vary over their likely range of values. More precisely, we assume that $U(k)$ is zero mean, and

$$E_q \begin{bmatrix} w(k-s) \\ U(k) \end{bmatrix} \begin{bmatrix} w^T(k-s), U^T(k) \end{bmatrix} = N_q N_q^T \quad (37)$$

where N_q is any square root of the covariance matrix above. As an example, if a feedback control of the form $u(k) = Gw(k)$ is used, then

$$U(k) = L_q w(k-s) \quad (38)$$

for a matrix L_q that is easily written in terms of G , A_q , B_q and M_q (but we omit the explicit details here), so that

$$N_q^T = [I \ L_q^T] \quad . \quad (39)$$

If process noise were also included, there would not be a deterministic coupling of $U(k)$ and $w(k-s)$, and a straightforward modification of (38) would provide the appropriate form for N_q .

Consider now the criterion (32), with all of the a_q taken to be 1 for the sake of simplicity. A direct calculation yields

$$J = \sum_{q=1}^Q \| P^T R_q \|^2_F \quad , \quad (40)$$

where

$$R_q = \begin{bmatrix} Z_q & H_q \\ 0 & I \end{bmatrix} N_q \quad , \quad (41)$$

so that the optimum choice of P is obtained from the singular value decomposition of $[R_1:R_2: \dots :R_Q]$.

4.3 Detection Versus Robustness

The methods described to this point involve measuring the quality of redundancy relations in terms of how small the resulting parity checks are under normal operating conditions. That is, good parity checks are maximally insensitive to modeling errors and noise. However, in some cases one might prefer to broaden the viewpoint. In particular, there may be parity checks that are not optimally robust (in the sense that we have discussed) but that are still of significant value because they are extremely sensitive to particular failure modes. In this subsection, we consider a criterion that takes such a possibility into account. We focus, for simplicity, on the noise-free case. The extension to include noise or known inputs as in the previous subsection is straightforward.

The specific problem to be considered is the choice of parity checks for the robust detection of a particular failure mode. We assume that the unfailed model of the system is

$$x(k+1) = A_q x(k) , \quad (42)$$

$$y(k) = C_q x(k) , \quad (43)$$

while if the failure has occurred the model is

$$x(k+1) = \bar{A}_q x(k), \quad (44)$$

$$y(k) = \bar{C}_q x(k) . \quad (45)$$

For example, if we return to the simple case $y_2(k) = ay_1(k)$, then under unfailed conditions one might have

$$a_1 \leq a \leq a_2 \quad (46)$$

while after a failure

$$\bar{a}_1 \leq a \leq \bar{a}_2 . \quad (47)$$

This is illustrated pictorially in Figure 3. In this case, one would like to choose the line P onto which one projects in such a way that a small projection is obtained if no failure has occurred and a large value results if a failure occurs. That is, we would like P to be "as orthogonal as possible" to Z and "as parallel as possible" to \bar{Z} .

Returning to the general problem, we again assume that q takes on one of Q possible values, and we let Z_q and \bar{Z}_q denote the counterparts of Z_q in (15) for the unfailed and failed models, respectively. We now have a tradeoff: we would like to make $P^T Z_q$ as small as possible for all q and to make $P^T \bar{Z}_q$ as large as possible. A natural criterion, for minimization over all P satisfying $P^T P = I$, is provided by

$$J = \sum_{q=1}^Q \left(\|P^T Z_q\|_F^2 - \|P^T \bar{Z}_q\|_F^2 \right) . \quad (48)$$

If we define the matrices

$$H = [Z_1:Z_2: \dots :Z_Q:\bar{Z}_1:\bar{Z}_2: \dots :\bar{Z}_Q] \quad (49)$$

and

$$S = \text{block diagonal } [I_{Qn} , -I_{Qn}] , \quad (50)$$




Figure 3: Illustrating Robust Detectability. Here \mathcal{Z} represents the set of values of (y_1, y_2) that can occur under normal operation, while $\bar{\mathcal{Z}}$ represents the corresponding set after the occurrence of a failure.

then

$$J = \text{tr} [P^T H S H^T P] \quad . \quad (51)$$

It is straightforward (see [3]) to show that a minor modification of the result in [17] leads to the following solution. We perform an eigenvector-eigenvalue analysis on the matrix

$$H S H^T = U \Lambda U^T \quad (52)$$

where U is orthogonal and

$$\Lambda = \text{diagonal} [\lambda_1, \dots, \lambda_N] \quad , \quad \lambda_1 \leq \dots \leq \lambda_N \quad . \quad (53)$$

Then the optimum choice for P is the first p columns of U :

$$P = [u_1 : \dots : u_p] \quad . \quad (54)$$

The corresponding minimum value of J in (48), (51) is

$$J^* = \sum_{i=1}^p \lambda_i \quad . \quad (55)$$

Two comments are in order about this solution. The first is that no more than Q_n of the λ_q can be positive. In fact the parity check based on u_q is likely to have larger values under failed rather than unfailed conditions if and only if $\lambda_q < 0$. Thus we immediately see that the maximum number of useful parity relations for detecting this particular failure mode equals the number of negative eigenvalues of $H S H^T$.

As a second comment, let us contrast the procedure we use here with the singular value decomposition of Z used in Section 3, which corresponds essentially to performing an eigenvector-eigenvalue analysis of ZZ^T . First, assume that precisely the first K of the λ_q are negative, and define

$$\begin{aligned} \sigma_1^2 &= -\lambda_1 \quad , \quad \dots \quad , \quad \sigma_K^2 = -\lambda_K \quad , \\ \sigma_{K+1}^2 &= \lambda_{K+1} \quad , \quad \dots \quad , \quad \sigma_N^2 = \lambda_N \quad , \end{aligned} \quad (56)$$

and

$$\Sigma = \text{diagonal } [\sigma_1, \dots, \sigma_N] . \quad (57)$$

From (52) we have that

$$HSH = U\Sigma S\Sigma^T U^T . \quad (58)$$

Assuming that Σ is nonsingular (which implies $K=Qn$), define

$$V = \Sigma^{-1} U^T H . \quad (59)$$

Then V is S-orthogonal,

$$VSV^T = S , \quad (60)$$

and H has what we call an S-singular value decomposition

$$H = U\Sigma V . \quad (61)$$

Thus, instead of the singular value decomposition of Z that we used in Section 3, the modified problem considered in this subsection calls for the S-singular value decomposition of H .

5. Discussion

This paper has developed methods for determining robust parity relations for failure detection in dynamic systems. The methods build on the geometric interpretation of parity checks as orthogonal projections of windows of observations onto subspaces that are as orthogonal as possible to the observation sequence, given the presence of model uncertainties and noise. We also considered the modification of this criterion to enable choice of parity checks for the detection of a particular failure mode. In each of the cases considered, a single singular value decomposition (or a variation of it, in the case of Section 4.3) produced a complete sequence of orthogonal parity relations, ordered in terms of a meaningful measure of robustness. In this section we provide brief discussions of several issues concerned with the interpretation and use of these results.

5.1 A Graphical Picture of Robust Redundancy

In all three of the formulations considered (in Sections 3, 4.1, and 4.2), we considered the problem of finding the p best parity checks. An obvious question, then, is what is a good value of p ? While our results do not give a precise answer to this question, they do provide a basis for obtaining a picture of the level of robust redundancy in a particular system configuration, as outlined next.

Recall that the solutions to our problems provide rank-ordered lists of parity relations, with a figure of merit for each relation given by a corresponding singular value (or eigenvalue for the case of Section 4.3). For example, consider the criterion (22). As we have seen, minimization of J over all choices of the parity check matrix P subject to the constraint that $P^T P = I_p$ (i.e. that we specify exactly p parity checks) results in the value J^{*p} given in (23), namely the sum of the p smallest singular values of the matrix Z in (18). The solid curve in Figure 4 illustrates a plot of this minimum value J^* as a function of p . Note that this curve must be convex, since the increment in J^* when we increase the number of parity checks from p to $p+1$ is σ_{p+1}^2 , which is at least as large as the squares of any of the p previous singular values. Furthermore, in this illustration the knee in the solid curve indicates a sharp increase in the singular values, which in turn points to a value of p beyond which the level of robustness decreases markedly.

Plots as in Figure 4 can also be of value in comparing different system configurations. In particular, in specifying a sensor complement




Figure 4: Illustrating the plot of the optimum value of the robustness criterion as a function of the number of parity checks specified (p takes on only integer values, but we have used continuous curves to facilitate illustration).

for a particular system, one is certainly interested in finding a set of sensors that provides a sufficient level of robust redundancy to allow accurate failure detection to be performed. Returning to Figure 4, the dashed line might correspond to the robust redundancy curve for an alternate sensor set. This set has a higher level of robust redundancy than the one corresponding to the solid line, since the dashed curve lies below the solid one. Clearly this is not a sufficient reason to state that the alternate sensor set is superior to the original one -- e.g. if the alternate set was obtained by adding several sensors to the original set, one would have to check that there is enough additional redundancy to permit the detection of the larger set of possible failures associated with this expanded sensor set -- but it does provide useful information for this design process.

Finally, we note that throughout the paper we have assumed a fixed order s for the parity checks under consideration. In any application one would, of course, want to consider several values of s . There are clear advantages (in terms of response time, and complexity of implementation) in considering small values of s , but the dynamics of a system may be such that there are important relationships of particularly high order. What one can imagine doing is solving the robust redundancy problem for $s = 1, 2, \dots$. Each such problem would result in a curve as in Figure 4, with the curve for each successive value of s lying below the preceding one. While this would appear to indicate that larger values of s always produce additional, useful parity checks, this is not necessarily the case -- one must check to see if these additional redundancy relations are truly useful or are simply nonminimal realizations of lower-order parity checks. For example, if $y_2(k) = ay_1(k)$, then $y_2(k) - ay_1(k)$ is a valid parity check, but so is $y_2(k) + y_2(k-1) - ay_1(k) - ay_1(k-1)$. See [3] for a polynomial matrix characterization of a complete set of minimal-order parity checks for deterministic linear systems and for a numerical example illustrating the issues raised in this section.

5.2 Alternate Robustness Criteria

In [2], Chow and Willsky consider a somewhat different formulation of the robust parity check problem. The criterion in [2] has several significant differences from the one we have used here, and in this section we describe the relationship between these. In the process we provide additional motivation for the present formulation. We also indicate several other criteria that in a sense represent intermediate steps between [2] and the present paper, and that provide some useful

insights. A more thorough development of these can be found in [3].

The model considered in [2] is a modified version of (27), (28) that includes known inputs and noise, and in which the model uncertainties are not constrained to a finite set of values. As discussed in Section 4.2 and the Appendix, there are direct ways in which one can incorporate known inputs and continuous parameter variations into the present formulation. The critical difference between [2] and our approach is the specific criterion chosen to define robustness. In particular, the principal problem posed and solved in [2] is the determination of the single best parity check $r(k)$ (so $p=1$), where "best" is defined as that with the minimum worst-case mean-squared value over the specified range of parameter uncertainties, with the system at a specified operating point -- i.e. the known input is assumed to take on a specified constant value, and the state $x(k-s)$ at the start of the data window is assumed to be at the equilibrium state corresponding to the constant control. While the consideration of operation at a particular set point does allow one to consider adapting parity checks to changing operating conditions, this flexibility is achieved at the expense of requiring that one solve a complex nonlinear optimization problem. Moreover, if one wishes to consider finding several parity checks, one must either solve one nonlinear optimization problem of greater complexity or a sequence of problems of equal complexity for each additional parity check.

As discussed in [3], if one removes the operating point constraint of [2] and assumes instead that the initial state is completely unconstrained, one is led to a criterion in which a parity space P has to be chosen to maximize either the minimum or average angle P makes with the observation space Z_q as q ranges over its full set of values. Here the cosine of the angle between two subspaces is defined as the maximum length of the projection of a unit vector from one space onto the other. While for any two subspaces this angle can be calculated using singular values [3], the maximization of the average or worst-case value of this angle is still a very complex nonlinear optimization problem. However, on reversing the steps of computing angles and averaging over parameter uncertainties, we are led to first compute a subspace that is the average of the Z_q and then choose P to be orthogonal to this average. This is very nearly the criterion we introduced in Section 3.

Specifically, as shown in [3] and [16], in this case we again choose the matrix Z_0 to minimize (17), but now with the columns of the matrices Z_q chosen to form orthonormal bases for the Z_q . The only

difference between this and the criterion we have used is the introduction in our case of scaling -- i.e. instead of viewing the initial state as completely unconstrained, we specify its covariance. With this specification we loose the interpretation of maximizing an angle between subspaces (since we replace orthonormal bases for the Z_q with the columns of the Z_q matrices defined in (15)), but the use of scaling is critical in order to obtain a practically meaningful criterion.

5.3 The Interpretation and Use of Parity Checks

Once we have determined a parity check, the question arises as to how this relation should be used. Chow and Willsky provide a detailed discussion of this issue in [2], and we shall not repeat it here. However, we shall make several brief comments in order to point to interesting avenues for further work.

Recall that the type of criterion on which we have focused in this paper is $E[\|r(k)\|^2]$, where the expectation is averaged over model uncertainty, noise, inputs, and initial conditions. This criterion is directly related to the performance of an open-loop [2] failure detection system in which the values of $r(k)$ calculated over an interval are used to make failure detection decisions (e.g. by comparing the sum of the squared norms of the $r(k)$ over the interval to a threshold).

It is also possible to use a parity check to define a closed-loop [2] failure detection algorithm. Specifically, as mentioned in Section 2, a parity check can be interpreted as defining a dynamic model. For example, a parity check of the form

$$r(k) = y_1(k) - y_1(k-1) - Ty_2(k-1) \quad (62)$$

(which might represent the relationship between the change in measured velocity, y_1 , to the measured acceleration, y_2 , scaled by the sampling time) can be interpreted as defining a model of the form

$$z_1(k) = z_1(k-1) + Ty_2(k-1) + w(k) \quad (63)$$

$$y_1(k) = z_1(k) + v_1(k) \quad (64)$$

where $z_1(k)$ represents the ideal noise-free value of y_1 , and the process noise $w(k)$ models both the expected deviations of $r(k)$ from zero under noise-free conditions (e.g. due to modeling error) and the presence of

sensor noise in $y_2(k-1)$. The model (63), (64) could then be used with any of the many existing sophisticated failure detection methods.

For example, one could consider basing failure detection decisions on the innovations $\mathcal{V}(k)$ from a Kalman filter based on (63), (64).⁺ A natural measure of robustness in this case would then be $E[\|\mathcal{V}(k)\|^2]$. This in turn raises the question of determining parity relations (i.e. finding P) to directly minimize $E[\|\mathcal{V}(k)\|^2]$. While this is an interesting and meaningful criterion, it is also true that this quantity is an extremely complicated and nonlinear function of P . Thus the methods of this paper would not directly apply to this problem, and it remains to be determined (a) if an efficient method can be obtained for solving this problem and (b) under what conditions, if any, significant performance improvements can be obtained by direct optimization of closed-loop innovations.

As a final comment, we note that the interpretation of parity checks as reduced-order models raises the question of whether the constructions developed here provide a useful, new method for model reduction. The exploration of this question remains for the future, but we note one interesting point. Specifically, what a parity relation such as (62) specifies is a constraint among the time evolutions of the components of $y(k)$. If one wishes to interpret such a relation as a dynamic model for the evolution of one of these components, as in (63), (64), then the other components of the measurement vector act as inputs to this model.

⁺ It is interesting to note that all but one of the parity relations used in [5,13] were used in an open-loop fashion. The remaining parity relation was used to design a second-order Kalman filter whose innovations were used to detect altimeter failures.

References

1. A.S. Willsky, "A Survey of Design Methods for Failure Detection," Automatica, Vol. 12, 1976, pp. 601-611.
2. E.Y. Chow and A.S. Willsky, "Analytical Redundancy and the Design of Robust Failure Detection Systems," IEEE Trans. Auto. Control, Vol. AC-29, No. 7, July 1984, pp. 603-615.
3. X. -C. Lou, "A System Failure Detection Method: The Failure Projection Method," S.M. Thesis, M.I.T. Dept. of Elec. Eng. and Comp. Sci.; also Rept. LIDS-TH-1203, M.I.T. Lab. for Inf. and Dec. Sys., June 1982.
4. R. Isermann, "Process Fault Detection Based on Modeling and Estimation Methods — A Survey," Automatica, Vol. 20, No. 4, 1984, pp. 387-404.
5. J.C. Deckert, M.N. Desai, J.J. Deyst, and A.S. Willsky, "F8-DFBW Sensor Failure Detection and Identification Using Analytic Redundancy," IEEE Trans. Auto. Control, Vol. AC-22, No. 5 Oct. 1977, pp. 795-803.
6. G.G. Leininger, "Sensor Failure Detection with Model Mismatch," 1981 Joint Automatic Control Conf., Charlottesville, Virginia, June 1981.
7. S.R. Hall, "Parity Vector Compensation for FDI," S.M. Thesis, Dept. of Aero. and Astro., M.I.T., Feb. 1981.
8. J.E. Potter and M.C. Suman, "Thresholdless Redundancy Management with Arrays of Skewed Instruments," in AGARDOGRAPH-224 on Integrity in Electronic Flight Control Systems, 1977, pp. 15-1 to 15-25.
9. M. Desai and A. Ray, "A Failure Detection and Isolation Methodology," 20th IEEE Conf. on Decision and Control, San Diego, California, Dec. 1981.
10. F.C. Schweppe, Uncertain Dynamic Systems, Prentice-Hall, Inc., Englewood Cliffs, N.J., 1973.
11. K.C. Daly, E. Gai, and J.V. Harrison, "Test for FDI in Redundant Sensor Configurations," AIAA J. Guidance and Control, Vol. 2, No. 1, Jan. 1979, pp. 9-17.

12. A.S. Willsky, E.Y. Chow, S.B. Gershwin, C.S. Greene, P.K. Houpt, and A.L. Kurkjian, "Dynamic Model-Based Techniques for the Detection of Incidents on Freeways," IEEE Trans. Auto. Control, Vol. AC-25, No., 3, June 1980, pp. 347-360.
13. J.C. Deckert, "Flight Test Results for the F-8 Digital Fly-by-Wire Aircraft Control Sensor Analytic Redundancy Management Technique," Proc. AIAA Guidance and Control Conference, Albuquerque, New Mexico, Aug. 1981, pp. 491-499.
14. V.C. Klema and A.J. Laub, "The Singular Value Decomposition: Its Computation and Some Applications," IEEE Trans. Aut. Control, Vol. AC-25, No. 2, April 1980, pp. 164-176.
15. G.H. Golub and C.F. Van Loan, Matrix Computations, The Johns Hopkins University Press, Baltimore, MD, 1983.
16. A. Bjorck and G.H. Golub, "Numerical Methods for Computing Angles Between Linear Subspaces," Math. Comput., Vol. 27, 1973, pp. 579-594.
17. C. Eckart and G. Young, "The Approximation of One Matrix by Another of Lower Rank," Psychometrika, Vol. 1, 1936, pp. 211-218.
18. D. W. Tufts, R. Kumaresan, and I. Kirsteins, "Data Adaptive Signal Estimation by Singular Value Decomposition of a Data Matrix," Proc IEEE, Vol. 70, No. 6, June 1983, pp. 684-685.

Appendix

Consider the problem of choosing an $N \times p$ matrix P to minimize

$$J = \sum_{q=1}^Q \|P^T Z_q\|_F^2 \quad (A.1)$$

subject to the constraint that $P^T P = I$. Note first that

$$J = \|P^T Z\|_F^2 = \text{tr}(P^T Z Z^T P) \quad (A.2)$$

where Z is defined in (16). As discussed in Section 3, we assume without loss of generality that $N \leq Qn$. Let the singular value decomposition of Z be as given in (18), (19).

We now show that the minimum value of J is

$$J = \sum_{i=1}^p \sigma_i^2 \quad (A.3)$$

and the optimum choice of P is

$$P = [u_1 : u_2 : \dots : u_p] \quad (A.4)$$

where the u_i are the first p left singular vectors of Z . To do this, we use the following elementary result, which is a direct consequence of the Courant-Fischer minimax principle [3, 14]: Suppose that

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \quad (A.5)$$

is $n \times n$, symmetric, and positive semidefinite. Suppose also that A_{11} is $m \times m$, and let $\lambda_i(A)$, $\lambda_i(A_{11})$ denote the i -th smallest eigenvalue of A , A_{11} respectively. Then

$$\lambda_i(A) \leq \lambda_i(A_{11}), \quad i = 1, \dots, m. \quad (A.6)$$

Consider then any choice of P satisfying the constraint $P^T P = I$, and augment this matrix with $N-p$ additional columns so that the square matrix

$$F = [P:D] \quad (A.7)$$

is orthogonal. Then

$$F^T Z Z^T F = \begin{bmatrix} P^T Z Z^T P & * \\ * & * \end{bmatrix}. \quad (A.8)$$

Applying (A.6) to (A.8) and using both (A.2) and the fact that F is orthogonal, we see that

$$\sum_{i=1}^P \sigma_i^2 = \sum_{i=1}^P \lambda_i(Z Z^T) = \sum_{i=1}^P \lambda_i(F^T Z Z^T F) \leq \text{tr}(P^T Z Z^T P) = \|P^T Z\|_F^2 \quad (A.9)$$

From (18) we see that

$$Z Z^T = U \Sigma U^T \quad (A.10)$$

with

$$\Sigma^T = \text{diagonal } [\sigma_1^2, \dots, \sigma_N^2] \quad (A.11)$$

From this we see that the inequality in (A.9) becomes an equality if p is chosen as in (A.4), thereby proving our assertion.

We note that from this analysis we can directly deduce that the same choice of p minimizes a variety of other criteria. For example, an interesting one is

$$\det(P^T Z Z^T P) \quad (A.12)$$

which has the interpretation of minimizing the volume of the projection of the columns of Z onto the subspace P. The proof that the same P minimizes (A.12) is also a straightforward consequence of (A.6) and (A.8). Specifically

$$\det(P^T Z Z^T P) = \prod_{i=1}^P \lambda_i(P^T Z Z^T P) \geq \prod_{i=1}^P \lambda_i(Z Z^T) = \prod_{i=1}^P \sigma_i^2 \quad (A.13)$$

with equality resulting once again if P is taken as in (A.4).

Finally, note that (as can be seen in (A.10)) we are actually using

the eigenvalue-eigenvector decomposition of

$$ZZ^T = \sum_{q=1}^Q Z_q Z_q^T$$

in order to find the optimal choice of P . This suggests a direct generalization of the criterion (A.1) to allow continuous parameter variations. Specifically, assume that $q \in K$, a compact subset of a finite-dimensional Euclidean space, and consider the following criterion:

$$J = \int_K \|P^T Z_q\|_F^2 dq = \text{tr}\{P^T (\int_K Z_q Z_q^T dq) P\} \quad (\text{A.14})$$

(As before, this can be interpreted as $E[\|r(k)\|^2]$, where we have absorbed the square root of the probability density of q into the definition of Z_q).

Consider the eigenvalue-eigenvector representation

$$\int_K Z_q Z_q^T dq = U \Lambda^T U \quad (\text{A.15})$$

where $0 \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$. Then the first p columns of U define the optimal choice of P . Note also that (assuming that $\lambda_1 > 0$) if we define

$$V_q = \Lambda^{-1/2} U^T Z_q \quad (\text{A.16})$$

then

$$Z_q = U \Lambda^{1/2} V_q \quad (\text{A.17})$$

where $U^T U = I$ and

$$\int_K V_q V_q^T dq = I \quad (\text{A.18})$$

Hence (A.17) is the singular value decomposition of the map Z_q .