SOME ASPECTS OF REAL-TIME MODEL VALIDATION TECHNIQUES FOR USE IN INTEGRATED SYSTEMS

P.J.G. Teunissen
Delft Geodetic Computing Centre (LGR)
Department of Geodesy, Delft University of Technology
Thijsseweg 11, 2629 JA Delft, The Netherlands

Abstract

The Kalman filter produces optimal estimators, but the quality of the estimators is only guaranteed of course as long as the assumptions underlying the mathematical model hold. Misspecifications in the model will invalidate the results of filtering and thus also any conclusion based on them. It is therefore of importance to have ways to verify the validity of the assumed mathematical model. In this paper a brief review is given of a general real-time detection, identification and adaptation (DIA) procedure for use in integrated navigation systems.

I. INTRODUCTION

It is well-known that the standard real-time minimum mean squared error navigation filter produces optimal estimators with well defined statistical properties [1]. The estimators are unbiased and they have minimum variance within the class of linear unbiased estimators. The quality of the estimators is however only guaranteed as long as the assumptions underlying the mathematical model hold. Misspecifications in the model will invalidate the results of filtering and thus also any conclusion based on them. It is therefore of importance to have ways to verify the validity of the assumed mathematical model.

The objective of the present paper is to give a brief review of the detection, identification and adaptation (DIA) procedure, that has been developed during the last few years at the Delft Geodetic Computing Centre [2-4]. It has been developed for the real-time model validation of integrated navigation systems and can be seen as the natural extension of the static *quality control theory* as described in [5-7]. The present review is partly based on the course "Special studies in numerical methods in geodesy and related surveying sciences" which the author was invited to give at the Department of Surveying of the University of Calgary in March 1990. Derivations of results are avoided in the paper. Instead we keep to the basic ideas involved and try to motivate the main results by appealing to intuition and geometric interpretations. Some variations and extensions of the theory, which are important for certain applications are mentioned but not further dwelt upon. For these more advanced procedures the reader is refered to the referenced literature.

The DIA-procedure discussed in the present paper consists of the following three steps

1. *Detection:* An overall model test is performed to diagnose whether unspecified model errors have occurred.

2. *Identification:* After detection of a model error, identification of the potential source of the model error is needed. This implies a search among the candidate hypotheses for the most likely alternative hypothesis and their most likely time of occurrence.

3. *Adaptation:* After identification of an alternative hypothesis, adaptation of the recursive navigation filter is needed to eliminate the presence of state vector biases.

The DIA-procedure is completely recursive and avoids the explicit use of a parallel bank of augmented navigation filters. It is based on the concept of a uniformly-most-powerful-invariant test-statistic and is applicable in principle to any dynamic system that fits into the frame work of the state-space formalism.

In: Proceedings of the AIG-Symposium on Kinematic Systems in Geodesy, Surveying and Remote Sensing, Banff, Canada, 10-13 September 1990. pp. 191-200.

II. FILTERING AND BIAS ACCUMULATION

It will be assumed that the reader is familiar with the standard navigation filter in state-space formalism [1]. Based on the usual model assumptions, the optimal recursive prediction and filtering equations for the state estimator read

$$\hat{\underline{x}}_{k,k-1} = \Phi_{k,k-1}\hat{\underline{x}}_{k-1|k-1} + \underline{d}_k \ , \ \hat{\underline{x}}_{k|k} = \hat{\underline{x}}_{k|k-1} + K_k[\underline{y}_k - A_k\hat{\underline{x}}_{k|k-1}] \ , \tag{1}$$

with corresponding variance matrices

$$P_{k|k-1} = \Phi_{k,k-1}P_{k-1|k-1}\Phi^*_{k,k-1} + Q_k \ , \ P_{k|k} = [I - K_kA_k]P_{k|k-1}. \tag{2}$$

This filter can be shown to produce optimal estimators of the state vector with well defined statistical properties. The state estimators are unbiased, are Gaussian distributed and have minimum variance within the class of linear unbiased estimators. It is important to realize however, that optimality is only quaranteed as long as the assumptions underlying the mathematical model hold. Misspecifications in the model will invalidate the results of estimation and this also any conclusion based on them. It is therefore of importance to have ways to verify the validity of the working hypothesis H_o for which (1) and (2) are optimal.

In order to verify the optimality of (1) and (2), the working hypothesis H_o will be opposed to a class of alternative hypotheses H_a. The specification of appropriate alternative hypotheses for a particular application is non-trivial and probably the most difficult task in the process of quality control. It depends to a great extent on experience and ones knowledge of the dynamic system. For the present discussion we restrict attention to misspecifications in the mean and assume that all second moments of the various random vectors are specified correctly and known. This restriction still leaves room for a sufficiently large class of alternative hypotheses that contains the most frequently occurring model errors. An extension of the theory for the case the second moments are only partially known is discussed in [8,9].

For a misspecification in the mean of the state vector we consider the following class of alternative hypotheses:

$$H_a : E\{\underline{d}_i\} = x_i - \Phi_{i,i-1}x_{i-1} - C_{x,i}\nabla_x \ , \tag{3}$$

where $E\{.\}$ is the mathematical expectation operator. The matrix $C_{x,i}$ is assumed to be known and the vector ∇_x is assumed to be unknown. Furthermore, it is assumed that $C_{x,i} \neq 0$ for $l \leq i \leq m$ and zero otherwise. The times l and m are unknown. Thus we assume to know the type of model error that may occur, but not the time period in which it occurs.

It will be clear that the slip $C_{x,i}\nabla_x$ in the dynamic model can typically accommodate under-parametrizations in the state vector. Assume for instance that the dynamics of a moving vehicle is based on a constant velocity model under H_o. Then, if at time l the vehicle starts accelerating linearly the constant velocity model becomes inadequate and an additional parametrization in the form of $C_{x,i}\nabla_x$ is needed. In a way that is completely analogous to (3), one can also model under the alternative hypothesis H_a, slips in the measurement model. A slip $C_{y,i}\nabla_y$ in the measurement model can typically accommodate outliers in the data, sensor failures and instrumental biases. For instance, an outlier at time l in the j-th observable can be modelled as $C_{y,l} = (0...10...0)^*$ with the 1 on the j-th place and $C_{y,i} = 0$ otherwise.

If H_a is true, filtering under H_o will generally result in biased estimators. It is therefore of importance to know how particular misspecifications in H_o manifest themselves as biases in the state vector or functions thereof. Knowledge of the impact of model errors can then be used to set acceptance criteria for the sizes of these model errors. This is of importance for the design of an appropriate navigation filter and for the design of a powerful enough DIA-procedure. Before considering the impact of model errors, one should first make clear on what functions of the state vector the impact is studied. This depends very much on the particular application for which the navigation filter is designed and it may range from just one single function of the state vector to all n elements of the state vector. For instance, the impact on instrumental parameters may or may not be of interest, or, one may be particularly interested in position but not in velocity, or, as is the case in some real-time GPS-applications, it is the horizontal solution which is of interest and not the pseudo-range bias. In all these cases one generally has to specify a set of linear(ized) functions of the state vector that is of particular interest. If we assume that these functions are collected in a matrix F^*, then it is the bias in $F^*\hat{\underline{x}}_{k|k}$ which is of interest. The bias in can be computed once a measure for ∇_x is available. An appropriate measure for ∇_x is provided by the *minimal detectable bias* (MDB). The MDB is defined as the model error ∇_x that can just be detected with the DIA-procedure at a fixed probability level. For more details on the concept of the MDB's and their relevance as a diagnostic tool, the reader is referred to [3,10].

With a measure for ∇_x available, the corresponding bias in $F^*\hat{\underline{x}}_{k|k}$ can be shown to follow as

$$F^*\nabla\hat{x}_{k|k} = F^*\Phi_{k|k+1}[\sum_{j=l}^{k}(\prod_{i=j}^{k}\Phi_{i+1,i}[I-K_iA_i])C_{x,j}]\nabla_x \tag{4}$$

The computation of (4) is straightforward and can be done *recursively* in a way that closely follows the recursion of the actual navigation filter. Once the bias of (4) is known, its significance can be tested with the following *bias-to-noise ratio*:

$$\lambda_{F^*\hat{x}_{k|k}} = \nabla\hat{x}^*_{k|k}F[F^*P_{k|k}F^*]^{-1}F^*\nabla\hat{x}_{k|k}\ . \tag{5}$$

Note that this scalar measure is *invariant* against reparametrizations of the state vector. Also note that (5) can be interpreted geometrically as the square of the length of the vector that follows from an orthogonal projection of $\nabla\hat{\underline{x}}_{k|k}$ onto the orthogonal complement of the nullspace of F^* (see figure 1).

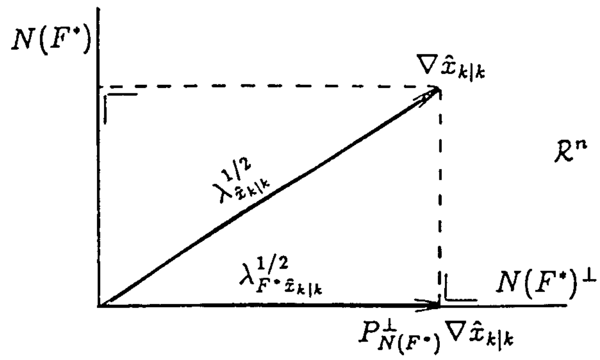

Fig. 1. $P^\perp_{N(F^*)}\nabla\hat{x}_{k|k}$ is the orthogonal projection of $\nabla\hat{x}_{k|k}$ along the nullspace of $F^*, N(F^*)$.

It can also be shown that (5) provides an *upperbound* for the bias-to-noise ratios of the individual elements of $F^*\hat{\underline{x}}_{k|k}$, see [3] but also figure 1. In the DIA-theory, the bias-to-noise ratio (5) is used together with the MDB's to set acceptance criteria for the impact of model errors. This can be done at the designing stage of the navigation filter.

If the bias-to-noise ratio (5) turns out to be unacceptable large and increasing as time proceeds, the usefulness of the actual navigation filter is of course completely nullified if the alternative hypothesis H_a is indeed the true hypothesis. This phenomenon is related to what is known in the Automatic Control literature as the *divergence* of the filter [11,12]: after a certain period of operation of the filter, the biases in the state estimators eventually diverge to values entirely out of proportions to the precision values predicted by the filter. This is illustrated for an elementary example in figure 2a. .

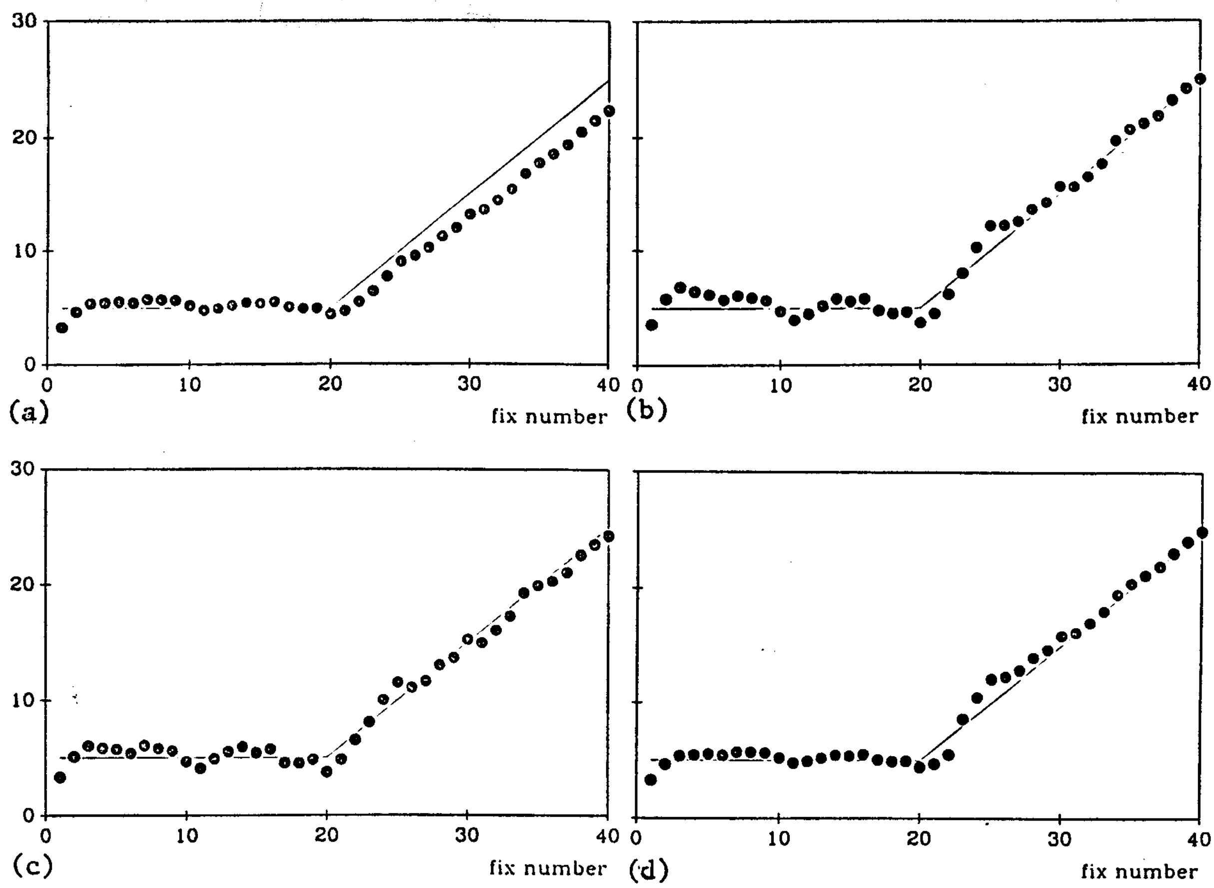

Fig. 2. (a) Divergence after fix 20, (b) Filtering under H_a for $k \geq 0$, (c) Filtering with increased process noise, (d) DIA at fix 23.

In this example, $\Phi_{k,k-1} = 1$ and $A_k = 1$ with $C_{x,i} = i - 20, i > 20$ and zero otherwise. There are a number of empirical techniques available that lesson such degradations due to model errors. One, albeit drastic approach is to include all potential sources of model errors in the mathematical model. This at least assures that the filtered state vectors are unbiased. But the disadvantage of this approach is of course that it leads to overparametrizations for those periods of time where H_o holds. This is shown in figure 2b, where filtering is done under H_a for the complete time period. A more subtle approach is based on increasing the process noise of the dynamic model. By increasing the process noise, one in effect increases the variance matrix of the predicted state as it is used in the filter. This has as consequence that less weight is given to the predicted state and that the bias accumulation gets damped. This is shown in figure 2c. The advantage of this approach is its conceptual simplicity and ease of implementation. And it has proven its value for a number of important applications [13]. The method has however also a number of drawbacks. First of all, one may question whether it is appropriate to use *second* moments as a tool to control the *first* moments. Of course, the increase in process noise may bound the impact of model errors and thus eliminate divergence. But it will never eliminate the presence of bias completely. A second drawback of increasing the process noise is that one obtains a filtered state estimator that is noiser than it needs to be for those time periods where H_o holds. This would not be the case, if one would know at what time instant to increase the process noise. But knowing this, implies knowing the starting time l of the model error. This leads therefore to the necessity of having available a *detector* of model errors. Finally, a third drawback of increasing the process noise is that it in effect reduces the *redundancy* *information* of the filter. Assume for instance that each measurement-update of the filter is based on only one measurement. A possible outlier in this measurement can then be detected by comparing the measured value with its predicted value based on the predicted state vector. The power of the outlier detection deteriorates however with an increase in the variancematrix of the predicted state. As a consequence, outliers in the data may pass unnoticed due to an increase in process noise and leave possible unacceptable biases in the filtered state vector.

A third approach of lessening degradations due to model errors is based on the concept of limited memory filtering [14,15]. The idea is here to base the filtering only on recent data in a moving window or fading window. In this approach the weight given to past information decreases as time proceeds. As a result the dependence of the present filtered estimator on past information decreases, with the effect that model errors are also damped out. The concept of limited memory filtering is closely related to the concept of increasing of the process noise. In fact, it can be shown that exponentially weighting of data is fully equivalent to an assumption of increased process noise [15].

Although the above approaches for eliminating divergence have proven their value for many important applications, it is the author's opinion that they are not really suited for handling alternative hypotheses as specified by (3). The methods are too empirical with no clearcut optimality properties. Nevertheless, the basic ideas of these approaches are sound and will (perhaps ironically) be seen to reappear in the DIA-theory.

III. DIA: DETECTION, IDENTIFICATION, ADAPTATION

Detection

The objective of the detectionstep is to test the *overall validity* of the mathematical model H_o. Detection is only possible if the navigation system provides for redundant information. It is this surplus of information which enables one to test the statistical consistency of the data with the model. Fortunately, the mathematical model on which the navigation filter is based has a built in redundancy, because of the presence of the dynamic model. This implies that the redundancy of each measurement-update k is equal to the number of measurements, say m_k, at this update. The statistical consistency checking of the data with the model is based on the *predicted residuals*. The predicted residual at time k, $\underline{v}_k$, is defined as the difference between the actual m_k-vector of observables at time k and the predicted vector of observables based on the predicted state vector: $\underline{v}_k = \underline{y}_k - A_k \hat{\underline{x}}_{k|k-1}$. It can be shown that under the working hypothesis H_o, the sequence of predicted residuals constitute a Gaussian distributed white noise process [4]. It is exactly this knowledge of the distribution of the predicted residuals under H_o that enables one to test the validity of the assumed mathematical model. Our *global overall model* (GOM) teststatistic for testing at time k, the overall validity of the mathematical model H_o is given as

$$\underline{T}^{l,k} = \underline{T}^{l,k-1} + g_k[\underline{T}^k - m_k \underline{T}^{l,k-1}], \tag{6}$$

with the input $\underline{T}^k = \underline{v}_k^* Q_{v_k}^{-1} \underline{v}_k$, where Q_{v_k} is the variancematrix of $\underline{v}_k$, and gain $g_k = 1/[\sum_{i=l}^{k} m_i]$ where m_i is the number of observables at time i (see figure 3).

The filter is initialized with $\underline{T}^{l,l} = 0$. Note that the GOM-teststatistic is computed recursively and that it is based on the predicted residuals which are readily available during each measurement update of the navigation filter. The GOM-teststatistic has been normalized, primarily for graphical purposes, such that its expectation under H_o is given as $E\{\underline{T}^{l,k}|H_o\} = 1$. Since its distribution under H_o follows a central F-distribution, the overall validity of H_o is rejected at time k and an unspecified model error is considered present in the time interval $[l, k]$ if and only if $T^{l,k} \geq F_\alpha(g_k, \infty, 0)$, where α is the upper probability point of the central F-distribution with g_k, ∞ degrees of freedom.

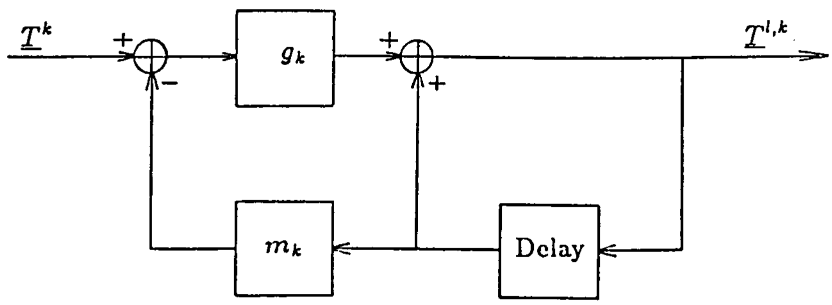

Fig. 3. The recursive GOM-teststatistic.

An important practical problem with the above GOM-test is the choice for l, the time that the model error is assumed to be starting to occur. Since the starting time of the model error is unknown a priori one has to start in principle with $l = 1$. But a fixed value for l, turns the filter (6) into a *growing-memory* filter, with the potential practical problem of a possibly long delay in time of detection. Rejection of H_o at time k, may imply namely that a global model error started to occur as early as time $l = 1$. In order to reduce the time of delay, a moving *window* of length N is introduced by constraining l to $k-N+1 \leq l \leq k$. When choosing N, one of course has to make sure that the detection power of the GOM-test is still sufficient. This is typical a problem one should take into account when designing the navigation filter. With the finite window of length N, the filter (6) is essentially reduced to a *finite-memory* filter.

Instead of using a finite window, one may alternatively use a fading window. By setting l equal to 1, and replacing the gain $g_k = 1/[\sum_{i=l}^{k} m_i]$ by the gain $g_k = w^k/[\sum_{i=l}^{k} m_i w^i]$, with $w \geq 1$ the filter reduces to a *fading-memory* filter. Note that $E\{\underline{T}^{l,k}|H_o\} = 1$ still holds with the fading window. The GOM-teststatistic however, instead of following an F-distribution now follows a linear combination of independent χ^2-distributions. With the fading window, the same recursive filter structure (6) is retained. This becomes advantageous when compared to the finite window, if a particular application requires the use of long windows. The weight factor w, which determines the nominal length of the fading window, is chosen on the basis of the detection power of the GOM-test. A useful approximation to the nominal length of the fading window is given by $(w-1)^{-1}$. A value of $w = 1.2$ would then correspond to a window length of $N = 5$. Although the type of window to use depends very much on the particular application at hand, one should keep in mind that the choice of the window length must always be based on the required detection power of the GOM-test.

The teststatistic (6) is termed the global overall model teststatistic, since it is designed to test the overall validity of the model and to detect unspecified global unmodelled trends. It can be shown that this teststatistic has the optimality property of being a uniformly-most-powerful-invariant teststatistic. Loosely speaking this means that one has with the GOM-test, the highest probability of correctly detecting unspecified model errors. In our applications of the DIA-theory, a distinction is made for practical reasons between local validity and global validity of the model. This distinction is introduced in order to have better detection and separation capabilities for model errors that have either a local or more global character. This implies that in the actual implementation of the theory detection consists of both a global overall model test and a local overall model test [2,4,9].

Identification

The next step after detection is the identification of the most likely model error. For identification, candidate alternative hypotheses need to be specified explicitly. For the present discussion we restrict attention to the class of alternative hypotheses as specified by (3). As with detection, identification gratefully makes use of knowledge of the distributional properties of the predicted residuals. In order to obtain a most sensitive teststatistic for identification purposes, we first need to know how a slip $C_{x,i}\nabla_x$ propagates into the predicted residuals. Let us denote the propagation of $C_{x,i}$, for $l \leq i \leq k$, into the space of predicted residuals by the matrix C_{v_k}. The matrix C_{v_k} can then be found from simply following the recursion of the navigation filter. This is shown in figure 4.

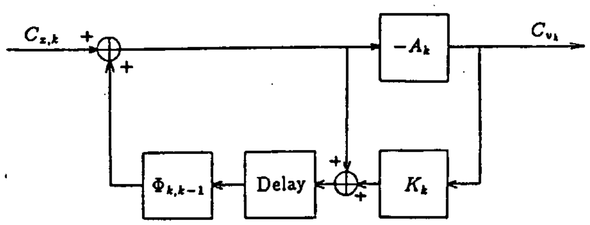

Fig. 4. The recursion of C_{v_k}.

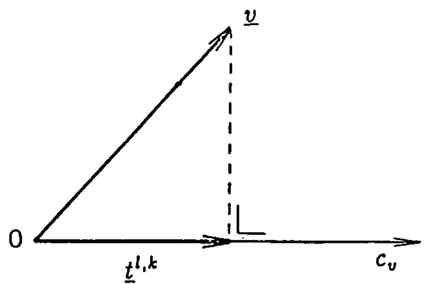

Fig. 5. The orthogonal projection: $\underline{t}^{l,k} = \dfrac{c_v^* Q_v^{-1}\underline{v}}{(c_v^* Q_v^{-1} c_v)^{1/2}}$.

With the matrix C_{v_k} available, we are now in the position to formulate the appropriate teststatistic. Let us assume for simplicity that the model error is one-dimensional and therefore that C_{v_k} is an m_k-vector, which will be denoted by the lower case kernel letter c_{v_k}. The appropriate slippage teststatistic for the identification of the slip $c_{x,i}\nabla_x, l \leq i \leq k$, follows then from an *orthogonal projection* of the vector of predicted residuals $\underline{v}$ onto the vector c_v. This is shown in figure 5. The corresponding *global slippage* (GS) teststatistic reads therefore

$$\underline{t}^{l,k} = \frac{\sum_{i=l}^{k} c_{v_i}^* Q_{v_i}^{-1} \underline{v}_i}{[\sum_{i=l}^{k} c_{v_i}^* Q_{v_i}^{-1} c_{v_i}]^{1/2}} \tag{7}$$

As will be intuitively clear, it is the orthogonal projection that quarantees that the GS-teststatistic (7) is most sensitive for the model error $c_{x,i}\nabla_x$. Note that, with Q_{v_i} and $\underline{v}_i$ readily available from the navigation filter and the efficient recursion of c_{v_i}, the computation of the GS-teststatistic parallels that of the actual navigation filter.

Strictly speaking, the GS-teststatistic (7), has to be computed for each alternative hypothesis considered and for each $k \geq l$. However, since l is unknown a priori, one has to start in principle with $l = 1$. This implies that one has to compute k number of teststatistics per alternative hypothesis at the time of testing k. As a result one obtains a *testmatrix* of increasing order with the GS-teststatistics as entries. This is shown in figure 6a. Clearly, this is unpractical, both from a computational point of view as well as because of the possible increase of the delay in time of identification. Fortunately, not all entries of the testmatrix may be necessary if one studies the power of the teststatistics. Although the power will increase theoretically for an increasing size of the interval $[l, k]$, the gain could be negligible for all practical purposes. This motivates, in accordance with our discussion of detection, the use of a moving window. This is shown in figure 6b for the case $k-N+1 \leq l \leq k$ and in figure 6c for the case $k-N+1 \leq l \leq k-M$. The rationale behind this last constraint is that in some applications the GS-teststatistic may be too insensitive for global model errors if $l > k-M$. Instead of using a finite window, the concept of a fading window, as discussed for detection can be applied to the teststatistic (7) as well.

$$
\begin{array}{cc} & \overrightarrow{k} \\ l\downarrow & \begin{bmatrix} t^{1,1} & t^{1,2} & t^{1,3} & t^{1,4} \\ & t^{2,2} & t^{2,3} & t^{2,4} \\ & & t^{3,3} & t^{3,4} \\ & & & t^{4,4} \end{bmatrix}_{H_a} \end{array}
\qquad
\begin{bmatrix} t^{1,1} & t^{1,2} & \bullet & \bullet \\ & t^{2,2} & t^{2,3} & \bullet \\ & & t^{3,3} & t^{3,4} \\ & & & t^{4,4} \end{bmatrix}_{H_a}
\qquad
\begin{bmatrix} \bullet & t^{1,2} & \bullet & \bullet \\ & \bullet & t^{2,3} & \bullet \\ & & \bullet & t^{3,4} \\ & & & \bullet \end{bmatrix}_{H_a}
$$

Fig. 6. Testmatrix of GS-teststatistic with (a) no window, (b) a moving window with $k-2 < l \leq k$ and (c) a moving window with $k-2 < l \leq k-1$.

With the windows introduced, we are now in the position to describe our identification procedure. At the time of testing k, one first determines per alternative hypothesis the value of l in the window for which the sample value of $\underline{t}^{l,k}$ is at a maximum. In other words, the kth column of the testmatrix is searched for the entry which is at a maximum. The corresponding row number of the testmatrix then identifies l as the most likely time of occurrence of the model error if the corresponding alternative hypothesis would be true. In order to find both the most likely alternative hypothesis and most likely value of l, the sample values of $max_{k-N+1 \leq l \leq k-M} \underline{t}^{l,k}$ for the different alternative hypotheses are compared. The maximum of this set finally identifies the most likely time of occurrence l and most likely alternative hypothesis.

The procedure described above contains the minimum requirements for identification. For some applications it may be necessary to develop a more advanced identification procedure. For instance, it may be necessary to discriminate between local and global slippages in the mean, or, to take care of possible masking effects due to different model errors, or to identify potential sources of model errors for which ∇_x is known. For the changes and/or extentions of the identification procedure which are required to accomodate these situations, the reader is referred to [4,8,9].

Adaptation

After identification of the most likely alternative hypothesis, adaptation of the recursive navigation filter is needed to eliminate the presence of biases in the filtered state of the navigation system. In order to be able to adapt the filter, one first needs an estimate of the identified model error ∇_x. It will be intuitivily clear that the estimate of ∇_x has to depend in some way on the predicted residuals. In fact one can show that the best estimator (in the sense of minimal variance) of the model error in the space of predicted residuals, is given by the orthogonal projection of the predicted residuals onto the range-space spanned by the columns of the matrix formed by the C_{v_i} [3]. With this result and the whiteness property of the predicted residuals follows then that the best estimator of ∇_x can be computed in a recursive form as shown in figure 7. This recursive filter is initialized with $\hat{\underline{\nabla}}^{l,l} = [C^*_{v_l} Q^{-1}_{v_l} C_{v_l}]^{-1} [C_{v_l} Q^{-1}_{v_l} \underline{v}_l]$ and its gainmatrix is given as $G_k = Q_{\hat{\nabla}^{l,k-1}} C^*_{v_k} [Q_{v_k} + C_{v_k} Q_{\hat{\nabla}^{l,k-1}} C^*_{v_k}]^{-1}$, where $Q_{\hat{\nabla}^{l,k-1}}$ is the variancematrix of $\hat{\underline{\nabla}}^{l,k-1}$.

With the estimator $\hat{\underline{\nabla}}^{l,k}$ available, we are now in the position to correct the filtered statevector at time k for the presence of bias. The adaptation consists of subtracting the *bias accumulation* due to $\hat{\underline{\nabla}}^{l,k}$ from the state filtered under H_o. The adapted filtered statevector reads

$$
\hat{\underline{x}}^a_{k|k} = \hat{\underline{x}}^o_{k|k} - X^l_{k,k} \hat{\underline{\nabla}}^{l,k}. \tag{8}
$$

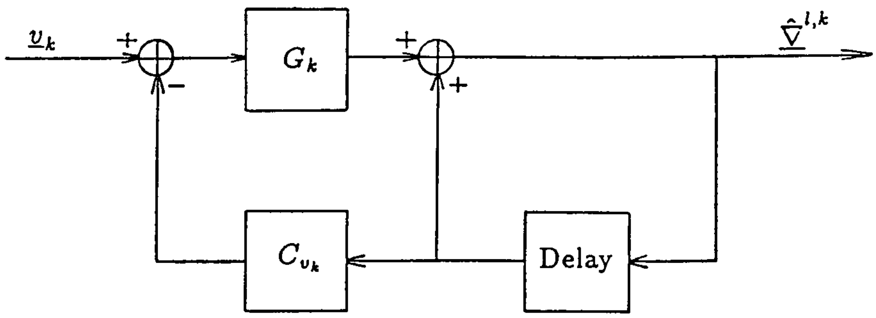

Fig. 7. The Recursion of $\hat{\underline{\nabla}}^{l,k}$.

The estimator $\hat{\underline{x}}^a_{k|k}$ is unbiased and corresponds with H_a. This shows that explicit filtering under H_a is not necessary and that the results can be based on the nominal navigation filter. The variance matrix of $\hat{\underline{x}}^a_{k|k}$ follows from applying the error propagation law to (8). Since $\hat{\underline{x}}^o_{k|k}$ and $\hat{\underline{\nabla}}^{l,k}$ can be shown to be uncorrelated, the variance matrix of $\hat{\underline{x}}^a_{k|k}$ follows simply as

$$P^a_{k|k} = P^o_{k|k} + X^l_{k,k} Q_{\hat{\underline{\nabla}}^{l,k}} X^{l*}_{k,k}. \tag{9}$$

The matrix $X^l_{k,k}$ of (8) and (9) corresponds with the bias accumulation matrix of (4). Both $X^l_{k,k}$ and $\hat{\underline{\nabla}}^{l,k}$ of (8) can be computed efficiently in recursive form. This follows essentially from a combination of figure 4 and figure 7. The corresponding recursion is shown in figure 8. The recursion of $X^l_{k,k}$ is initialized with $X^l_{l,l} = [I - K_l A_l] C_{x,l}$.

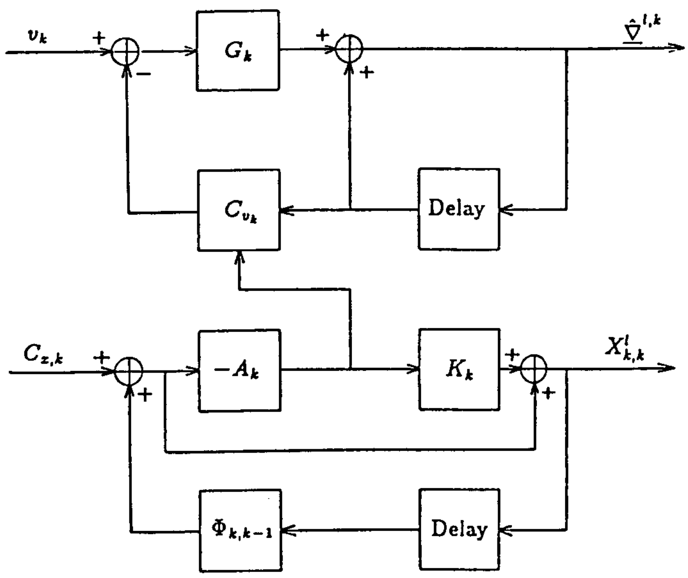

Fig. 8. The recursion of $\hat{\nabla}^{l,k}$ and $X^l_{k,k}$.

With (8), adaptation takes place at time k. This implies that strictly speaking the filtered states remain biased between l and k. Thus it would be theoretically more correct to adapt all filtered states in the time interval $[l, k]$. This is possible and may even be required for some applications. The corresponding adaptation involves smoothing, but is again based on an expression like (8). The reason why we have focussed in the present discussion on the above simple approach of only adapting $\hat{x}^o_{k|k}$, is twofold. Firstly, in real-time applications it is the present estimate of the state vector which is of interest and not so much the past estimators. And secondly, the bias in the state vector for times between l and k may considered to be negligible if the built up of the model error is still too small to be detected with the GOM-teststatistic. Note that this again points out, that one should choose the window length and detection power in relation to the biases one is willing to accept.

References

[1] Gelb, A. (Ed.) (1974): Applied Optimal Estimation. The M.I.T. Press.

[2] Teunissen, P.J.G. and M.A. Salzmann (1988): Performance Analysis of Kalman Filters. Proceedings HYDRO88, Special Publication, No. 23, Hydrographic Society, pp. 185-193.

[3] Teunissen, P.J.G. (1989): Quality Control in Integrated Navigation Systems. IEEE Aerospace and Electronic Systems Magazine, Vol. 5, No. 7, pp. 35-41.

[4] Teunissen, P.J.G. and M.A. Salzmann (1989): A Recursive Slippage Test for Use in State-Space Filtering. Manuscripta Geodaetica, Vol. 14, No. 6, pp. 383-390.

[5] Baarda, W. (1968): A Testing Procedure for Use in Geodetic Networks. Netherlands Geodetic Commision, Publications on Geodesy, New series, Vol. 2, No. 5, Delft.

[6] Teunissen, P.J.G. (1984): Quality Control in Geodetic Networks. In: Optimization and Design of Geodetic Networks, Grafarend, E. and F. Sansò (Eds.), Springer Verlag, Berlin, pp. 526-547.

[7] Staff of DGCC (1982): The Delft Approach for the Design and Computation of Geodetic Networks. In: Forty Years of Thought, Delft, pp. 202-274.

[8] Teunissen, P.J.G. (1986): Adjusting and Testing with the Models of the Affine and Similarity Transformation. Manuscripta Geodaetica, Vol. 11, pp. 214-225.

[9] Teunissen, P.J.G. (1990): An Integrity and Quality Control Procedure for Use in Multi-Sensor Integration. Paper presented at ION Satellite Divison, 3rd International Technical Meeting, Colorado Springs, USA.

[10] Salzmann, M.A. (1990): MDB: A Design Tool for Integrated Navigation Systems. Paper presented at Int. Symp. on Kinematic Systems in Geodesy, Surveying and Remote Sensing, Banff, Canada.

[11] Fitzgerald, R.J. (1971): Divergence of the Kalman Filter. IEEE Transactions on Automatic Control, Vol. AC-16, No. 6, pp. 736-747.

[12] Schlee, F.H., C.J. Standisch and N.F. Toda (1966): Divergence of the Kalman Filter. AIAA Journal, Vol. 5, No. 6, pp. 1114-1120.

[13] Jazwinski, A.H. (1970): Stochastic Processes and Filtering Theory. Academic Press, New York.

[14] Sorenson, H.W. and J.E. Sacks (1971): Recursive Fading Memory Filtering. Information Science, 3.

[15] Anderson, B.D.O. (1973): Exponential Data Weighting in the Kalman-Bucy Filter. Information Science, 5, pp. 217-230.