Detection of Interference/Jamming and Spoofing in a DGPS-Aided Inertial System

NATHAN A. WHITE
PETER S. MAYBECK, Fellow, IEEE
STEWART L. DeVILBISS
Air Force Institute of Technology

Previous research at the Air Force Institute of Technology (AFIT) has resulted in the design of a differential Global Positioning System (DGPS) aided INS-based (inertial navigation system) precision landing system (PLS) capable of meeting the FAA precision requirements for instrument landings. The susceptibility of DGPS transmissions to both intentional and nonintentional interference/jamming and spoofing must be addressed before DGPS may be safely used as a major component of such a critical navigational device. This research applies multiple model adaptive estimation (MMAE) techniques to the problem of detecting and identifying interference/jamming and spoofing in the DGPS signal. Such an MMAE is composed of a bank of parallel filters, each hypothesizing a different failure status, along with an evaluation of the current probability of each hypothesis being correct, to form a probability-weighted average state estimate as an output.

For interference/jamming degradation represented as increased measurement noise variance, simulation results show that, because of the good failure detection and isolation (FDI) performance using MMAE, the blended navigation performance is essentially that of a single extended Kalman filter (EKF) artificially informed of the actual interference noise variance. However, a standard MMAE is completely unable to detect spoofing failures (modeled as a bias or ramp offset signal directly added to the measurement). This work describes a *moving-bank pseudoresidual* MMAE (PRMMAE) to detect and identify such spoofing. Using the PRMMAE algorithm, spoofing is very effectively detected and isolated; the resulting navigation performance is equivalent to that of an EKF operating in an environment without spoofing.

Manuscript received February 28, 1997; revised November 21, 1997.

IEEE Log No. T-AES/34/4/07976.

Authors' address: Dept. of Electrical and Computer Engineering, Air Force Institute of Technology, 2950 P Street, Building 640, Wright-Patterson AFB, OH 45433-7765.

0018-9251/98/$10.00 1998 IEEE

I. INTRODUCTION

Currently, the Department of Defense and the commercial airline industry are using the Instrument Landing System (ILS) for aircraft guidance during precision approaches. Recent precision landing research conducted by the FAA and other researchers has established the Global Positioning System (GPS) as a primary sensor for a precision landing system (PLS) to replace the aging ILS [1–4]. Particularly applicable to this research is the work of Gray [2] (using standard GPS) and Britton [1] (using differential GPS (DGPS)), which has shown that an integrated GPS-aided inertial navigation system (INS) based PLS meets FAA requirements for Category I and II precision approaches.

The accuracy potential of GPS as a primary sensor in a PLS is well established. However, possible interference (benign or malignant) of the low-power GPS signal remains a major concern in such a safety-of-flight critical system. Integration of GPS with other navigation sensors (using an extended Kalman filter (EKF)) provides a possible solution to this interference problem (see [5–7] for a thorough presentation of Kalman filter theory and applications). This work discusses the use of multiple model adaptive estimation (MMAE) to detect and identify (and so compensate for) GPS failures in a DGPS-aided INS-based integrated navigation system.

II. PROBLEM STATEMENT

The integrated PLS used in this work is composed of an INS, barometric altimeter, radar altimeter, DGPS, and a ground-based GPS pseudolite (an additional GPS transmitter placed at a known location on the ground, rather than onboard a satellite). This integrated system has 64 truth model (i.e., best simulation model) states and 13 states in the filter model. The truth model for the integrated PLS is made up of 22 DGPS states, 39 INS states, two baro altimeter states, and one pseudolite state. The radar altimeter error is modeled by white noise and uses no states. The filter error model has two DGPS states, nine (Pinson [5: pp. 305, 343]) INS states, and two baro altimeter states. Since error state space modeling is used for the filters in this application, the measurements are actually the difference between sensor measurements external to the INS (differential range measurements to four satellites and to one pseudolite, and radar altimeter and barometric altimeter outputs) and the INS-indicated value of the same variables. Explicit modeling details are in [8]. With this given system configuration having demonstrably good performance attributes under nominal conditions [1, 2], the question is, can acceptable performance be maintained despite the interference/jamming or attempted spoofing of the

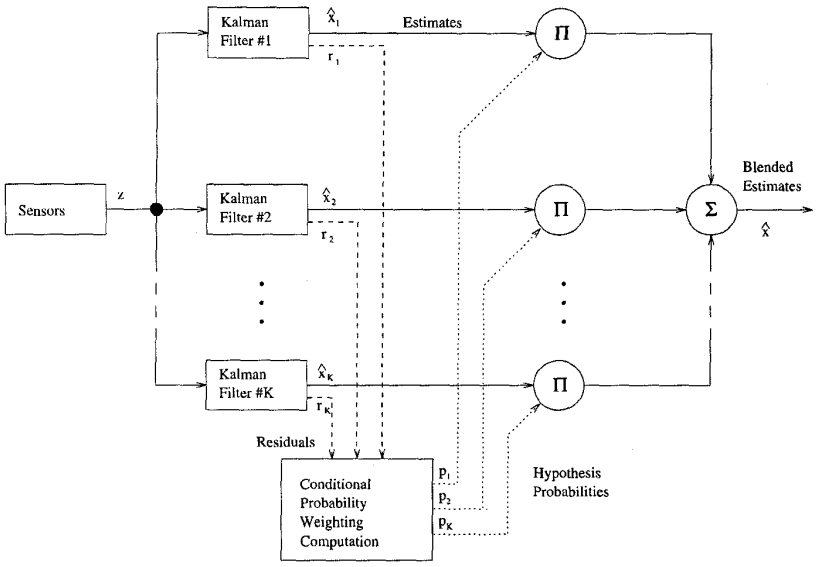

Fig. 1. Multiple model adaptive estimation (MMAE).

GPS, using only the information in the residuals of the systems-integration filter (i.e., without use of GPS receiver power measurements to detect jamming, or use of multiple GPS antennas to address spoofing, etc.)?

The goal of this work is to demonstrate the failure detection and isolation (FDI) capability that is available using MMAE in the integrated system, not only for military and commercial users, but potentially for civilian aircraft using very low-cost components. To this end, simulations were conducted using four different navigation component cases, the least precise and cheapest of which is composed only of a low-quality (4.0 nm/hr) baro-aided INS and DGPS. This work presents the results of only one navigation case, that of a high-quality (0.4 nm/hr) baro-inertial system combined with both DGPS (including pseudolite) and radar altimeter, but similar FDI results were obtained [8] for each of the four navigation cases used.

III. MMAE OVERVIEW

Fig. 1 shows a functional block diagram of the MMAE algorithm. Its primary feature is a bank of sampled-data Kalman filters operating in parallel, each using an identical measurement environment [6, 9, 10]. Each Kalman filter models the dynamics of the system (in this case the PLS) under different conditions of no-fail or failed operation, such as greatly increased measurement noise covariance $\mathbf{R}(t_i)$, which might represent the effects of GPS interference/jamming. At each sample period, each of these K filters produces a state estimate $\hat{\mathbf{x}}_k$, and a vector of residuals $\mathbf{r}_k$, for $k = 1,2,\ldots,K$. The residuals produced by a filter with an accurate model of the real world are zero-mean, white, Gaussian, and of covariance as computed within that filter itself. During operation, the filter having the most well-behaved residuals (to be made more precise subsequently) contains the model which best matches the true status of the system with regard to interference/jamming or spoofing [5, 6, 11–13].

Status identification takes place in the conditional hypothesis probability block. Based on the residuals, a probability of model adequacy p_k ranging from zero to one is computed for each elemental filter within the MMAE structure, according to the following recursive equation:

$$p_k(t_i) = \frac{f_{\mathbf{z}(t_i)|\mathbf{a},\mathbf{Z}(t_{i-1})}(\mathbf{z}_i \mid \mathbf{a}_k, \mathbf{Z}_{i-1}) p_k(t_{i-1})}{\sum_{j=1}^{K} f_{\mathbf{z}(t_i)|\mathbf{a},\mathbf{Z}(t_{i-1})}(\mathbf{z}_i \mid \mathbf{a}_j, \mathbf{Z}_{i-1}) p_j(t_{i-1})} \quad (1)$$

where

$$f_{\mathbf{z}(t_i)|\mathbf{a},\mathbf{Z}(t_{i-1})}(\mathbf{z}_i \mid \mathbf{a}_k, \mathbf{Z}_{i-1}) = \frac{1}{(2\pi)^{m/2}|\mathbf{A}_k(t_i)|^{1/2}} \exp\{\cdot\} \quad (2)$$

$$\{\cdot\} = \{-\tfrac{1}{2}\mathbf{r}_k^T(t_i)\mathbf{A}_k^{-1}(t_i)\mathbf{r}_k(t_i)\}.$$

In the above equations, $f_{\mathbf{z}(t_i)|\mathbf{a},\mathbf{Z}(t_{i-1})}(\mathbf{z}_i \mid \mathbf{a}_k, \mathbf{Z}_{i-1})$ is the probability density function of the current measurement $\mathbf{z}(t_i)$, of dimension m, conditioned on the hypothesized failure status $\mathbf{a} = \mathbf{a}_k$ and previously observed measurement history $\mathbf{Z}(t_{i-1})$, based on a filter's residuals $\mathbf{r}_k(t_i)$ and internally computed residual covariance $\mathbf{A}_k(t_i)$. The individual filter residual vectors are formed as the difference between the actual measurement when it arrives and the filter's best prediction of the measurement based on its assumed model of the real world. For the sampled-data EKF,

$$\mathbf{r}_k(t_i) = \{[\mathbf{z}(t_i) - \mathbf{h}_k[\hat{\mathbf{x}}_k(t_i^-), t_i]\} \quad (3)$$

where $\mathbf{h}_k$ is the measurement vector model in the kth elemental filter, and the residual covariance is computed by the kth elemental filter as

$$\mathbf{A}_k(t_i) = [\mathbf{H}_k(t_i)\mathbf{P}_k(t_i^-)\mathbf{H}_k^T(t_i) + \mathbf{R}_k(t_i)] \quad (4)$$

where $\mathbf{H}_k$ is the partial of $\mathbf{h}_k$ with respect to $\mathbf{x}$. When a filter's actual residuals are in consonance with the filter-computed covariance $\mathbf{A}_k$, the exponential term in (2) is approximately $-m/2$, where m is the measurement dimension. In filters modeling an incorrect hypothesis, the residuals are larger than anticipated by $\mathbf{A}_k$ and the magnitude of this negative term is greater, thus causing a deweighted p_k for that hypothesis. Because of the recursive nature of (1), in practice the p_ks are given an artificial lower bound $p_{k_{\min}} = 0.01$ to prevent "lock-out" (without the lower bound, if any p_k is ever computed as zero, it will be zero thereafter). The 0.01 bound was chosen as the largest value (allows fastest p_k growth) that preserved "good" operation in performance simulations, i.e., the Bayesian blended state estimates are essentially unaffected by the state estimates of the lower-bounded elemental filters. The vector of probabilities produced by the conditional hypothesis probability block is used to identify interference/spoofing status and to weight the state estimates of the individual filters. The output of this algorithm is a Bayesian blended estimate of the system states, i.e., the probability-weighted average $\hat{\mathbf{x}}$ as seen in Fig. 1. A Bayesian blended estimate of the system parameters $\mathbf{a}$ (failure status in this case) is similarly obtained by weighting each $\mathbf{a}_k$ with that hypothesis' corresponding p_k (for $k = 1,2,\ldots,K$, as shown in Fig. 1), and summing.

IV. MOVING-BANK MMAE

The MMAE algorithm propagates multiple filters forward in time, continually selecting the filter, or weighted combination of filters, that appears to have (on the basis of residual characteristics) the best model of the real world. Often, the possible parameter space is so large that completely discretizing it requires many more filter hypotheses than can realistically be run in parallel. In this case, a small group of filters, with assumed parameter values that are in the close neighborhood of the current parameter estimate, are chosen to be active at any given time. If the estimated parameter moves significantly, the bank of active filters is moved to be centered around the new parameter estimate. This algorithm is called a moving-bank MMAE.

The state dynamics of the PLS necessitate the use of a moving-bank MMAE for the accurate tracking of spoofing failures. As modeled in this research (see Section V), the effect of spoofing failures is equivalent to the effect of a change in GPS user clock bias. Because of this similarity, the GPS user clock bias filter state in a non-moving-bank MMAE quickly (during a single filter propagation cycle) assumes the spoofing-induced error, causing filter divergence due to the now-inconsistent state vector. Using a moving-bank MMAE with spoof offset (bias in indicated GPS differential range measurement) as the uncertain parameter, a detected spoofing event initiates a sub-algorithm in which the spoofing offset magnitude is identified and the bank is recentered on that new spoof offset value *before* the next filter propagation cycle is entered. In this way the current spoofing offset is "remembered" while the filter itself experiences no erroneous change in its state vector estimate.

V. FAILURE MODELS

Interference is modeled as a sudden increase in the measurement noise $\mathbf{R}$ associated with all GPS satellite vehicles (SVs). GPS *jamming* refers to the total loss of GPS transmissions. Such a failure is well and easily modeled as a *very large* increase in $\mathbf{R}$. This results in a very light weighting for those measurements in the elemental Kalman filters; the effect is essentially the same as if those measurements were never received.

Spoofing is modeled as a bias added to the measurements associated with all GPS SVs. The addition of this bias (if it is undetected) induces a bias on the GPS position solution. An intelligent spoofer would place unique offsets on each of the SV pseudo-range measurements, producing a specific position offset desired by the spoofer. In this research, however, spoofing biases are inserted as either step or ramp offsets added uniformly to all GPS SV measurements. The results obtained using the simplified model extend directly to the more complex spoofing model [8].

Specific magnitudes of interference or spoofing will not be considered; each failure was allowed to assume a range of values in order to demonstrate the capability of MMAE to detect and identify interference failures of varying magnitude, and of a moving-bank pseudoresidual MMAE (PRMMAE) to identify spoofs of various sizes and types.

VI. FILTER STRUCTURE

A three-filter MMAE bank was used to detect and isolate interference/jamming failures. The level of measurement noise for the DPGS measurements is given by $\mathbf{R}_0$, and it is assumed that $\mathbf{R} = 2000 \times \mathbf{R}_0$ is the maximum level of real-world interference which will be encountered. It has been experimentally determined that assumed values of $\mathbf{R}$ for the individual Kalman filters within the MMAE should be separated by about an order of magnitude in order to prevent identification ambiguity [8]. Thus, the three interference/jamming filter assumptions were $\mathbf{R} = \mathbf{R}_0$,

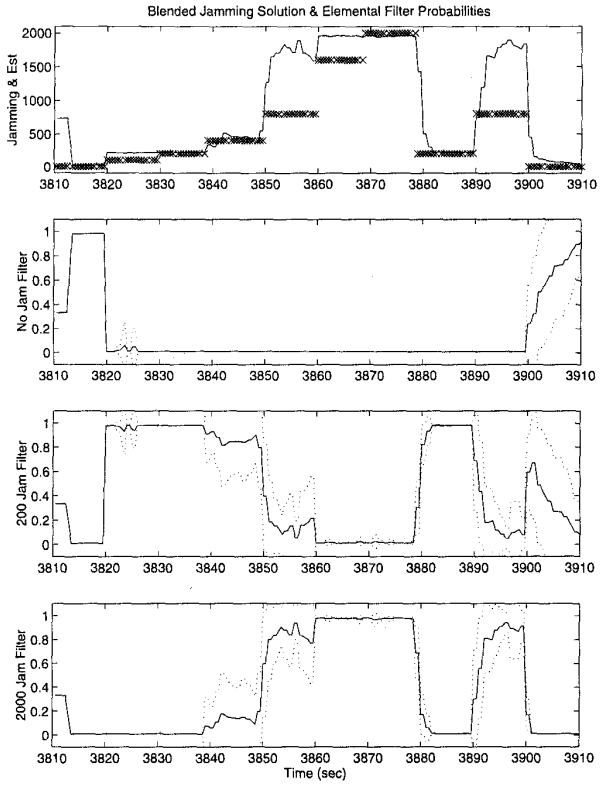

Fig. 2. Interference/jamming failures, three elemental filters. FDI performance (sub-plot 1) and elemental filter p_k mean $\pm$ std. dev. values (sub-plots 2–4).

$\mathbf{R} = 200 \times \mathbf{R}_0$, and $\mathbf{R} = 2000 \times \mathbf{R}_0$. This filter bank did not move within the parameter space.

The spoofing MMAE filter bank must be arranged in symmetric pairs about a no-bias filter in order to detect and isolate spoofing biases of either a negative or positive sign. The spoofing bias assumed by an individual filter is given by $\mathbf{b}_s$. One filter pair is located, in parameter space, quite closely to the no-bias filter, in order to provide good hypothesis blending to identify small spoofing changes or to recenter the filter bank after a large move. A second symmetric filter pair is located further away from the no-bias filter, near the experimentally determined limit of useful blending performance, in order to capture larger spoofing bias changes. Thus, the five-filter spoofing MMAE bank assumptions are $\mathbf{b}_s = 0$ ft (no bias, relative to the currently estimated spoof offset value), $\mathbf{b}_s = \pm 15$ ft, and $\mathbf{b}_s = \pm 240$ ft. Recall that the spoofing filter bank just described is required to remain centered on the actual bias in order to preserve good state estimation and future spoofing offset tracking.

VII. RESULTS

Figs. 2, 3, and 4 of this section are organized to show the FDI operation and performance of the MMAE/moving-bank PRMMAE algorithm clearly and concisely. The first subplot of these figures shows the FDI performance of the algorithm; the x's denote the true (simulated) value of the failure parameter, and the line is the average (over the 15 sample runs) estimate of the **R** multipliers (additive spoof offset in the case of spoofing failure Figs. 3 and 4). The remaining three subplots of Fig. 2 (the last five subplots for spoofing failure Figs. 3 and 4) show the time histories of mean (solid line) $\pm$ one standard deviation (dotted lines) p_k probability values for each active elemental filter. Subplot two of Fig. 3 shows the spoof offset estimation *error* mean (x's) $\pm$ one standard deviation (dotted). In Fig. 4, the spoof offset estimation error information is included in the first subplot. The next subplot in Figs. 3 and 4 shows the number of spoof estimation operations required to recenter the filter bank at each measurement time. The results discussed and plotted in this section are produced by 15-run Monte Carlo performance evaluations.

A. Interference/Jamming (Measurement Noise Failure) Results

Because the measurement noise **R** strongly influences the elemental filter probabilities p_k (see (1)–(4)), it was anticipated, before simulations were conducted, that the MMAE would be quite effective at

detecting and identifying, through Bayesian blending, interference/jamming failures. The **R** parameter estimation provided by the three-filter (assuming 1x, 200x, and 2000x the original tuned **R** values) interference bank, along with the p_k probabilities for each elemental filter, is shown in Fig. 2. As can be seen from the figure, the MMAE algorithm very quickly detects the onset of interference/jamming. The Bayesian blended state estimation performance which results from the MMAE FDI method is very good; in fact, it is essentially the same as that produced by a single EKF artificially informed of the actual measurement noise changes [8].

As expected, MMAE FDI was very effective against interference failures. However, as can be seen in Fig. 2, the filter bank has some difficulty in blending its two noise hypotheses. With the aid of hindsight, it could have been expected that the MMAE algorithm would be prone to bouncing probability weighting from hypothesis to hypothesis without much blending, and to have a tendency to choose the model with the larger assumed noise level. As the real world noise level becomes even slightly greater than that assumed by one filter, that filter's residuals very quickly look bad (i.e., its $[\mathbf{r}_k^T(t_i)\mathbf{A}_k^{-1}(t_i)\mathbf{r}_k(t_i)]$ becomes a large value) because the measurements are consistently violating its assumptions. Although the actual measurement noise variance is much less than it expects, the filter assuming a larger noise value sees measurements that do consistently fall within its expected variance (its $[\mathbf{r}_k^T(t_i)\mathbf{A}_k^{-1}(t_i)\mathbf{r}_k(t_i)]$ is considerably smaller). Probability quickly flows to the filter with the larger noise variance assumption.

B. Spoofing (Measurement Bias Failure) Results

In this section the filter is designed to detect *only* spoof failures, whereas only interference/jamming failures were discussed in the previous section. Because it is desired to have only a small number of elemental filters active at one time, it is expected that it will be required to implement a moving-bank MMAE in order to detect and identify arbitrarily sized spoofing (measurement bias) jumps. Early experimental trials revealed two important and related considerations affecting such an implementation. First, a correctly implemented moving-bank MMAE algorithm is *completely* unable to distinguish between the residuals of the individual elemental filters based on different presumed spoof offset values: no information is available on which to base the recentering of the filter bank. Second, this ambiguity among the filter probabilities precludes useful hypothesis blending even within the portion of parameter space spanned by the active filter bank. Both of these effects are products of the GPS user clock bias state's similarity to the induced spoofing failure: that state estimate in each elemental filter is corrupted by the difference between the actual spoof offset and the offset assumed by that filter.

A new technique, PRMMAE, provides useful information as to how to recenter the filter bank. Good parameter and state estimation is achieved and maintained by performing bank recentering operations before entering the next filter propagation cycle. This recentering prevents the central (no-bias) filter's state estimates from ever being corrupted by the spoof offset. The mathematical motivation and development of the PRMMAE algorithm is shown below.

At all times let the actual measurement be

$$\mathbf{z}_{\text{true}} = \mathbf{H}_{\text{true}}\mathbf{x}_{\text{true}} + \mathbf{v}_{\text{true}} + \mathbf{b}_{\text{true}}. \tag{5}$$

Let us assume two filter models, 1) assuming no measurement bias, i.e., that $\mathbf{b}_{\text{true}} = \mathbf{0}$, and 2) assuming a positive measurement bias, i.e., that $\mathbf{b}_{\text{true}} = \mathbf{b}_1$, where $\mathbf{b}_1$ will, at least initially, be assumed to be the bias actually simulated in the real world. These two filters will have, as their best prediction $\hat{\mathbf{z}}$ of the measurement before it arrives, $\hat{\mathbf{z}}_1 = \mathbf{H}\hat{\mathbf{x}}_1^-$ and $\hat{\mathbf{z}}_2 = \mathbf{H}\hat{\mathbf{x}}_2^- + \mathbf{b}_1$, which will produce the following update equations used to generate $\hat{\mathbf{x}}^+$ at each sample period:

$$\begin{aligned}\hat{\mathbf{x}}_1^+ &= \hat{\mathbf{x}}_1^- + \mathbf{K}[\mathbf{z}_{\text{true}} - \mathbf{H}\hat{\mathbf{x}}_1^-]\\ \hat{\mathbf{x}}_2^+ &= \hat{\mathbf{x}}_2^- + \mathbf{K}[\mathbf{z}_{\text{true}} - (\mathbf{H}\hat{\mathbf{x}}_2^- + \mathbf{b}_1)].\end{aligned} \tag{6}$$

Let $\mathbf{b}_{\text{true}} = \mathbf{0}$ for a period of time. After allowing both filters to run to steady state, each filter will modify its state estimates so that the residual vector tends to be zero-mean, i.e., $E[\mathbf{z}_{\text{true}} - \hat{\mathbf{z}}] = \mathbf{0}$. In fact, it has been observed that the elemental filters primarily altered their *user clock bias state* to yield such zero-mean residuals. This has the following effect:

$$\begin{aligned}E[\mathbf{H}\hat{\mathbf{x}}_1^-] &= \mathbf{z}_{\text{true}}(\mathbf{b}_{\text{true}} = \mathbf{0}) = \mathbf{H}_{\text{true}}\mathbf{x}_{\text{true}}\\ E[\mathbf{H}\hat{\mathbf{x}}_2^-] &= \mathbf{z}_{\text{true}}(\mathbf{b}_{\text{true}} = \mathbf{0}) - \mathbf{b}_1 = \mathbf{H}_{\text{true}}\mathbf{x}_{\text{true}} - \mathbf{b}_1.\end{aligned} \tag{7}$$

At the onset of a spoof, $\mathbf{b}_{\text{true}}$ becomes non-zero, $\mathbf{b}_{\text{true}} = \mathbf{b}_1$ for example. At that measurement update time, the *true* residuals are

$$\begin{aligned}[\mathbf{z}_{\text{true}} - \hat{\mathbf{z}}_1] &= [(\mathbf{H}_{\text{true}}\mathbf{x}_{\text{true}} + \mathbf{v}_{\text{true}} + \mathbf{b}_1) - \mathbf{H}\hat{\mathbf{x}}_1^-]\\ [\mathbf{z}_{\text{true}} - \hat{\mathbf{z}}_2] &= [(\mathbf{H}_{\text{true}}\mathbf{x}_{\text{true}} + \mathbf{v}_{\text{true}} + \mathbf{b}_1) - (\mathbf{H}\hat{\mathbf{x}}_2^- + \mathbf{b}_1)]\end{aligned} \tag{8}$$

whereas the *pseudo*residuals, namely $[\mathbf{z}_{\text{true}} - \mathbf{H}\hat{\mathbf{x}}_k^-]$ for all k, are

$$\begin{aligned}[\mathbf{z}_{\text{true}} - \mathbf{H}\hat{\mathbf{x}}_1^-] &= [(\mathbf{H}_{\text{true}}\mathbf{x}_{\text{true}} + \mathbf{v}_{\text{true}} + \mathbf{b}_1) - \mathbf{H}\hat{\mathbf{x}}_1^-]\\ [\mathbf{z}_{\text{true}} - \mathbf{H}\hat{\mathbf{x}}_2^-] &= [(\mathbf{H}_{\text{true}}\mathbf{x}_{\text{true}} + \mathbf{v}_{\text{true}} + \mathbf{b}_1) - \mathbf{H}\hat{\mathbf{x}}_2^-].\end{aligned} \tag{9}$$

The *true* residuals, using $E[\mathbf{H}\hat{\mathbf{x}}_k^-]$ from (7), both become $\mathbf{b}_1 + \mathbf{v}_{\text{true}}$, while the *pseudo*residuals, using $E[\mathbf{H}\hat{\mathbf{x}}_k^-]$ from (7), become 1) $\mathbf{b}_1 + \mathbf{v}_{\text{true}}$ and 2) $2\mathbf{b}_1 + \mathbf{v}_{\text{true}}$.

In writing (5)–(9), it is assumed that *true* residuals are used to update the elemental filters; what is subject to consideration here is whether the true residuals or the pseudoresiduals are more useful in forming $[\mathbf{r}_k^T \mathbf{A}_k^{-1} \mathbf{r}_k]$ for the MMAE's probability computations. As seen above, the true residuals from each elemental filter do not reveal the real world measurement bias and are indistinguishable from one another. It may be deduced that the filter assuming the *negative* of the actual bias will show nearly zero-mean *pseudo*residuals at the measurement update immediately following the spoof onset. Hence, good identification may be achieved by using the *pseudoresiduals* to form $[\mathbf{r}_k^T \mathbf{A}_k^{-1} \mathbf{r}_k]$. In the above discussion, had a filter assumed a bias of $-\mathbf{b}_1$, then in steady state its $E[\mathbf{H}\hat{\mathbf{x}}^-]$ (as in (7)) would have become $\mathbf{H}_{\text{true}}\mathbf{x}_{\text{true}} + \mathbf{b}_1$. The pseudoresiduals of such a filter would be roughly zero-mean at the measurement update following the spoof onset (and this information is only visible for a single measurement update time). The information provided by this zero-mean measurement pseudoresidual is used to isolate the actual bias. *The true residuals must be used (as in a single Kalman filter or regular MMAE) to update the elemental filters.*

Once the pseudoresidual information is gathered from the filters within the MMAE, the spoof is estimated (the value that would appear to be needed to zero the residual in the middle filter), and then that current measurement is *reprocessed* assuming that the estimated spoof is present in the real world (*the entire filter bank is moved to the neighborhood of the estimated spoof*). During operation, filter bank movement is accomplished by subtracting $\mathbf{b}_{\text{estimated}}$ from the true measurements before they enter the MMAE algorithm, rather than adding $\mathbf{b}_{\text{estimated}}$ to each filter's $\hat{\mathbf{z}}_k$, since it is computationally more efficient. This process is repeated for that single update time until the new spoof bias value is completely identified and the bank is recentered. Complete failure identification is assumed when each elemental filter p_k changes by less than 0.2 from one iteration to the next. After this correction, each of the elemental filters in the bank steps forward into the next propagation cycle without knowing it ever experienced a spoof. This process can be observed in the p_k plots of simulations involving spoofing elemental filters. In the first, third, and fifth subplots of Fig. 4 (to be seen and described in detail later in this section, associated with ramp rather than step spoofs, but called out here because it clearly shows the desired phenomenon), for example, it can be seen that when the real-world spoof bias changes by +15 ft/s (the first spoofing ramp), the p_k associated with the −15 ft bias elemental filter spikes for one second. The bank is moved based on this spike, and in the next second the probability returns to the no-bias elemental filter. Throughout the simulation depicted in Fig. 4 (and each spoofing simulation), it may be seen that the p_k value corresponding to the no-bias elemental filter displays such downward spikes each time a change in the real-world spoof bias value is detected. These downward spikes depict the information associated with the *first* measurement update which caused the MMAE filter bank to be moved; in reality, due to the measurement reprocessing described above, the no-bias elemental filter's p_k value never moves from its near-unity value.

Fig. 3 shows the detection and isolation performance of the PRMMAE algorithm in the moving-bank configuration against step-valued spoofs. Note several aspects of Fig. 3 which are in accordance with the description of the moving-bank PRMMAE algorithm's operation as described so far. It can be seen that the probability rests consistently (excepting the downward spikes mentioned above) on the no-bias elemental filter (fourth sub-plot), while the real-world and declared spoof values range over 2000 ft (first sub-plot). This is a clear indication that the no-fail filter is consistently recentered on the real-world spoof value. Each time the real-world spoof value jumps by a value modeled in one of the elemental filters, the probability spikes on the filter assuming the *negative* of the spoof jump value (sub-plots five through eight). The third sub-plot (measurement count) shows the number of measurement update iterations that were taken to identify the new spoof offset value fully at each sample time before the following propagation cycle was entered.

Fig. 3 indicates that spoofing jumps are identified accurately regardless of whether their value is exactly modeled in the bank of elemental filters. The spoof offset values are identified well enough that the resulting state estimation performance is, as in the case of interference/jamming failures, equivalent to that of a single Kalman filter artificially informed of the actual spoof offsets. The PRMMAE algorithm alone is unable to identify spoofing values greatly different from those modeled in its bank. The difficulty here is that the MMAE residual information term in (2), $\mathbf{r}_k^T \mathbf{A}_k^{-1} \mathbf{r}_k$, becomes very large when the spoof offset is numerically greater than 150 (ft) displaced from the bias assumption of a given elemental filter. Because $-1/2$ times this large term appears in the exponential of (2), the computed value of that equation (for use in (1)) goes to zero and the MMAE algorithm cannot make a decision about which direction to move the bank in the possible failure space. Even though the MMAE calculations using the $\mathbf{r}_k^T \mathbf{A}_k^{-1} \mathbf{r}_k$ term become useless, the individual measurement residuals associated with the DGPS measurements in each elemental filter allow simple isolation of the spoof magnitude under this condition. These residuals are zero-mean up until the addition of the spoof. When the spoof occurs, it shows up directly on all the residuals. Estimation is a matter of

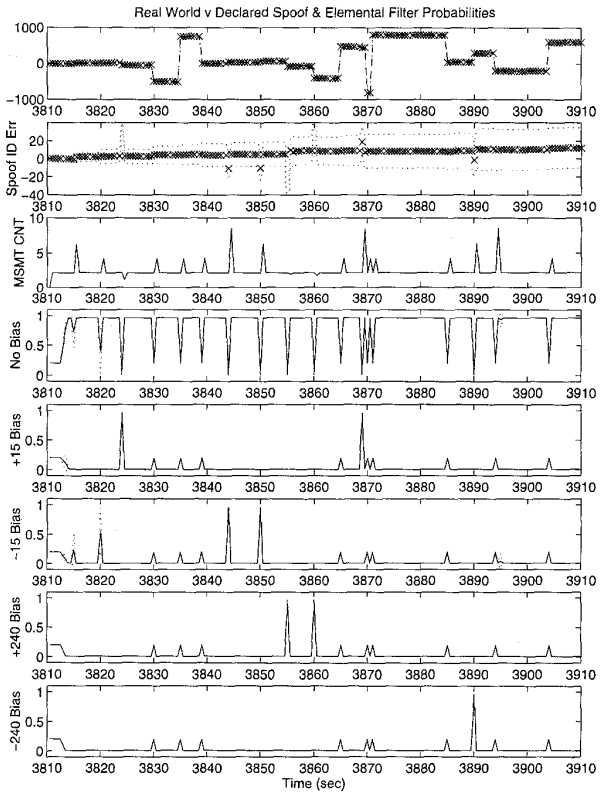

Fig. 3. Spoofing (steps) FDI performance and elemental filter probabilities: moving-bank PRMMAE.

simply reading the number. It can be seen in Fig. 3 (sub-plots 1, 5–8) that large changes in the real-world spoof value correspond to probability values of $0.2 = 1/(\#\text{ filters})$ for every filter. When all filters in the bank have equally bad residuals, the probability is equally divided between the filters and the individual residual terms are examined for the spoof offset as described above.

Identification of spoofing as described in the previous paragraph is accomplished equally well with only two bias-assuming filters symmetrically displaced about a filter assuming no bias offset, and not the four symmetrically displaced filters that were used earlier in this research. The three-filter spoofing bank can be used to identify the spoof onset and to fine-tune the spoof estimation once the first guess is made.

Fig. 4 shows the FDI performance of the moving-bank PRMMAE algorithm in the face of ramped (intelligent) spoofing offsets. The first ten real-world spoofing ramp segments (sub-plot one) have slopes of 30, 24, 20, 16, 12, 10, 8, 6, 4, and 2 ft/s, respectively. It can be seen that no significant identification error occurs until the 8 ft/s spoofing segment (note the dotted mean $\pm\,\sigma$ trajectories on the first subplot). The smaller spoofing ramps do cause considerable confusion to the algorithm, although the state estimation still does not degrade appreciably [8]. It can be concluded that, given the navigation component configuration used, spoofing steps and spoofing ramps as small as 10 ft/s can be identified with no significant error. When the actual spoofing error is smaller than this observable lower sensitivity bound, the PRMMAE algorithm is unable to distinguish the spoof offset (measurement bias) from the measurement noise clearly, and the spoof offset estimation error ceases to be zero-mean. While the spoof estimation error is still zero-mean, the resulting state estimation performance is equivalent to that of an artificially informed EKF.

An important consideration is inherent throughout the above discussion. The elemental filters of the moving-bank PRMMAE must be allowed to initialize with no measurement bias present in the real world. What is actually detected and tracked is the offset of the measurement bias from what was present during initialization. This assumption is reasonable because GPS initialization will presumably be accomplished before take-off from a friendly air base. Any attempted spoofing there would be detected by surveyed receivers at the station, and shortly removed.

VIII. CONCLUSIONS AND RECOMMENDATIONS

This paper shows the development of moving-bank PRMMAE, a new technique for the identification of

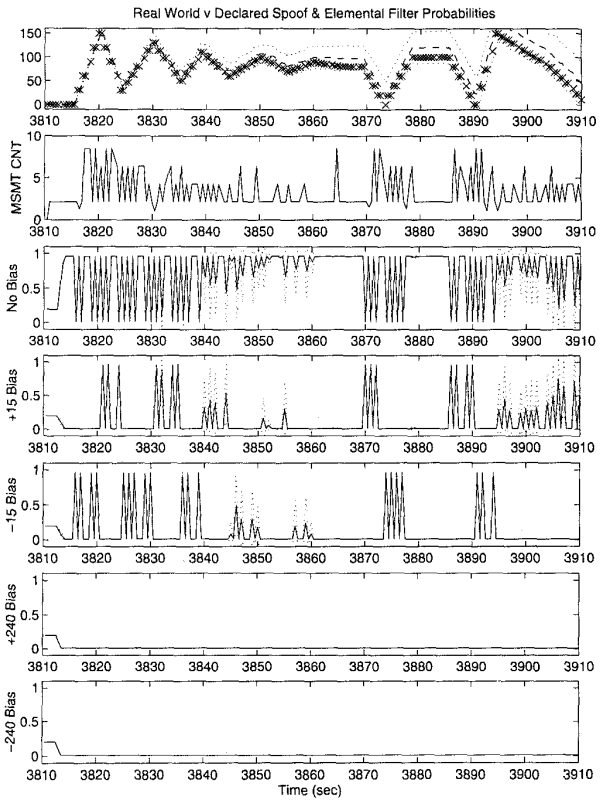

Fig. 4. Spoofing (ramps) FDI performance and elemental filter probabilities: moving-bank PRMMAE.

measurement offset (bias or ramp "spoof") events. Moving-bank PRMMAE (for spoofs) and standard MMAE (for interference) are used to detect and compensate for interference and spoofing events in the GPS portion of a GPS-aided INS navigation configuration. Accurate state estimation is maintained before, during, and after these events. These results indicate that a PLS based on these navigation components can reliably detect degradation of the navigation solution due to external RF sources, and can preserve the quality of navigation so that flight and some categories of instrument landings may safely be continued. It is expected that any system subjected to increased measurement noise or to measurement offset events will experience similar performance benefits through the application of the FDI techniques described in this paper.

For differenced measurements (DGPS), multipath, which may produce errors less than several meters, is the dominant residual error source. Multipath is a bias phenomenon (always a delay) which appears on the individual pseudo-range measurements. A subject of possible future application of the techniques developed in this paper is seen by noting the similarity between GPS multipath errors and the spoofing (measurement bias) errors addressed in this paper. The limiting factor in such an application is the sensitivity limit of moving-bank PRMMAE (see Section VIIB), which at approximately 10 ft (bias change per measurement update) excludes application to many instances of multipath. It was discovered however, that by artificially weighting the spoof bias estimates of the ± 15 symmetric filter pair, zero-mean spoof offset estimation was acheived for measurement bias steps as small as 1 ft. The artificial weighting is accomplished by implementing a p_k lower bound greater than is necessary just to prevent probability lock-out. In this case cited, a lower bound of 0.06 produced zero-mean estimation of measurement bias offsets of one foot, a factor of ten increase in sensitivity over the performance in Section VIIB associated with a lower bound of 0.01. While a complete mathematical explanation of this phenomenon has yet to be developed, such performance argues for the application of this method to detect and isolate multipath errors as well as the onset of spoofing.

REFERENCES

[1] Britton, R. L. (1995)
A differential GPS-aided INS for aircraft landings.
M.S. thesis, AFIT/GE/ENG/95D-03, School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH, Dec. 1995.

[2] Gray, R. A. (1994)
An integrated GPS/INS/BARO and radar altimeter system for aircraft precision approach landings.
M.S. thesis, AFIT/GE/ENG/94D-13, School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH, Dec. 1994.

[3] Rowson, S. V., et al. (1994)
Performance of category IIIB automatic landings using C/A code tracking differential GPS.
In *Proceedings of the Institute of Navigation 1994 National Technical Meeting*, San Diego, CA, 1994, 759–767.

[4] Wullschleger, V., et al. (1994)
FFA/Wilcox flight test results of DGPS system for precision approach.
In *Proceedings of the Institute of Navigation 49th Annual Meeting*, Cambridge, MA, 1994, 111–118.

[5] Maybeck, P. S. (1979)
Stochastic Models, Estimation, and Control, Vol. I.
New York: Academic Press, 1979. Republished, Arlington, VA: Navtech, 1994.

[6] Maybeck, P. S. (1982)
Stochastic Models, Estimation, and Control, Vol. II.
New York: Academic Press, 1982. Republished, Arlington, VA: Navtech, 1994.

[7] Maybeck, P. S. (1982)
Stochastic Models, Estimation, and Control, Vol. III.
New York: Academic Press, 1982.

[8] White, N. A. (1996)
MMAE detection of interference/jamming and spoofing in a DGPS-aided inertial system.
M.S. thesis, AFIT/GE/ENG/96D-21, School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH, Dec. 1996.

[9] Magill, D. T. (1965)
Optimal adaptive estimation of sampled stochastic processes.
IEEE Transactions on Automatic Control, **AC-10**, 5 (1965), 434–439.

[10] Athans, M., and Chang, C. B. (1976)
Adaptive estimation and parameter identification using a multiple model estimation algorithm.
Technical note 1976-28, ESD-TR-76-184, Lincoln Laboratory, MIT, Lexington, MA, June 1976.

[11] Van Trees, H. L. (1968)
Detection, Estimation and Modulation Theory.
New York: Wiley, 1968.

[12] Chow, E. Y., and Willsky, A. S. (1980)
Issues in the development of a general design algorithm for reliable failure detection.
In *Proceedings of the IEEE Conference on Decision and Control*, Albuquerque, NM, Dec. 1980, 1006–1012.

[13] Kalaith, T. (1970)
The innovations approach to detection and estimation theory.
Proceedings of the IEEE, **58** (1970), 680–695.

Nathan A. White was born December 21, 1970 in Providence, UT. He received a B.S. in electrical engineering from Utah State University, Logan, in 1995, and was commissioned into the United States Air Force. In Dec. 1996, he received an M.S. in electrical engineering from the Air Force Institute of Technology.

Following graduation, he was assigned to the Director's staff, Avionics Directorate of Wright Laboratories, Wright-Patterson AFB, OH.

Peter S. Maybeck (S'70—M'74—SM'84—F'87) was born in New York, NY on February 9, 1947. He received the B.S. and Ph.D. degrees in aeronautical and astronautical engineering from M.I.T., Cambridge, in 1968 and 1972, respectively.

In 1968, he was employed by the Apollo Digital Autopilot Group of The C. S. Draper Laboratory, Cambridge, MA. From 1972 to 1973, he served as a military control engineer for the Air Force Flight Dynamics Laboratory and then joined the faculty of the Air Force Institute of Technology in June of 1973. He is currently Professor of Electrical Engineering, responsible for the graduate sequence in estimation and stochastic control and for individual advanced digital filtering and control courses. His current research interests concentrate on using optimal estimation techniques for guidance systems, tracking, adaptive systems and failure detection purposes.

Dr. Maybeck is author of numerous papers on applied optimal filtering as well as the book, *Stochastic Models, Estimation and Control* (Academic Press, Vol. 1–1979, Vols. 2 and 3–1982; republished by Navtech in 1994). He is a member of Tau Beta Pi, Sigma Gamma Tau, Eta Kappa Nu, and Sigma Xi. He was recipient of the DeFlorez Award (ingenuity and competence of research), the James Means Prize (excellence in systems engineering) and the Hertz Foundation Fellowship at M.I.T. in 1968. In all years from 1975 to 1997, he received commendation as outstanding Professor of Electrical Engineering at A.F.I.T. In December of 1978, he received an award from the Affiliate Societies Council of Dayton as one of the twelve outstanding scientists in the Dayton, OH area. In March of 1980, he was presented with the Eta Kappa Nu Association's C. Holmes MacDonald Award, designating him as the outstanding electrical engineering professor in the United States under the age of 35 (he had placed second in this national competition for 1977 as well). In 1985, he received the Frederick Emmons Terman Award, the highest national award to a Professor of Electrical Engineering given by the American Society of Engineering Education. He is a member of the AIAA, and he is the current IEEE Dayton Section Student Activities Chairman and a member of the IEEE Executive Committee of Dayton, and he previously served as Chairman of the local Automatic Control Group.

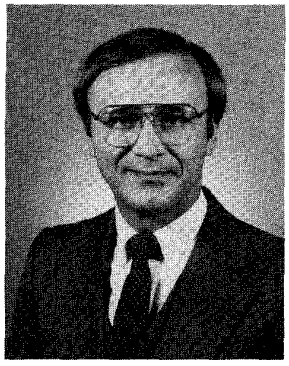

Stewart L. DeVilbiss received the B.S. degree in electrical engineering from the University of Missouri, Columbia, in 1986, the M.S. degree in electrical engineering from Purdue University, Lafayette, IN, in 1987, and the Ph.D. degree in electrical engineering from the Ohio State University, Columbus, in 1994.

He is a Captain in the U.S. Air Force and has previously served as a navigation engineer for the C-17 System Program Office and a flight control engineer for the Flight Dynamics Directorate of Wright Laboratory. He has been a member of the faculty of the Department of Electrical Engineering at the Air Force Institute of Technology since 1994. His current research area of interest is Global Positioning System technology and applications.