

Event Detection and Characterization During Autonomous Interplanetary Navigation

Timothy P. Crain*

NASA Johnson Space Center, Houston, Texas 77058

Robert H. Bishop†

University of Texas at Austin, Austin, Texas 78759-5321

and

Todd A. Ely‡

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109

A prototype autonomous/adaptive interplanetary navigation system employing neural networks and genetic algorithms is introduced. This system consists of a near real-time autonomous monitoring component and an offline adaptive modeling component. Unexpected dynamic or measurement events trigger 1) alerts to navigators and 2) recommendations for modifying the appropriate model parameters to improve filtering. The autonomous component analyzes tracking measurement residuals to first detect and then characterize the unexpected event. Once an environment change has been detected and characterized, the adaptive modeling component then modifies the necessary model parameters to bring the tracking filter back into optimal operation. The autonomous monitor employs a hierarchical mixture-of-experts model where the experts are extended Kalman filters organized into banks regulated by two levels of single-layer neural networks. The autonomous monitor is the focus of this presentation and demonstrates the ability to detect successfully the occurrence and to differentiate between the characteristics of unexpected discrete velocity changes and continuous dynamic changes. These environment changes are represented by an unmodeled impulsive maneuver and by solar radiation pressure surface mismodeling, respectively. The robust decision-making capability of this approach is further demonstrated by successfully characterizing three successive environment changes. All experiments were performed on recorded Mars Pathfinder two-way Doppler data from the period of 4 February 1997 to 17 April 1997.

Introduction

CURRENT tracking of interplanetary spacecraft relies heavily on navigator experience and ad hoc techniques for the resolution of anomalous measurement residual signatures produced by the operational tracking algorithm or filter. These anomalies are the result of changes in the dynamics and measurement system that cause the spacecraft environment to deviate from the tuned system model employed by the operational filter. Difficulties navigating recent interplanetary missions, such as modeling the solar radiation pressure (SRP) surface on Mars Pathfinder (MPF)¹ and modeling the thruster performance of Mars Climate Orbiter and Mars Polar Lander, illustrate the need for a systematic and expedient method for resolving anomalous behavior in the tracking solution. A prototype autonomous/adaptive navigation system employing neural networks and genetic algorithms has been developed. With this system, unexpected dynamic or measurement events trigger 1) alerts to navigators and 2) recommendations for modifying the appropriate model parameters to improve filtering. The architecture of this system consists of a near real-time autonomous monitoring component and an offline adaptive modeling component, as illustrated in Fig. 1.

The autonomous monitoring component operates in parallel with the real-time operational filter, processing measurements as they become available. Conceptually, the autonomous monitor could op-

erate either onboard or at a ground tracking station, but is assumed to operate at the tracking station here. Operational filter measurement residuals are compared to the residuals from a set of hypothetical filters contained in a hierarchical mixture-of-experts (HME) framework.^{2,3} Each of the filters in the HME is an expert on a specific mode of unmodeled environment change. If a transition to suboptimal modeling is detected, the nature of the mismodeling responsible for the transition is characterized by a reassignment in weights assigned by the HME. Once characterized, the adaptive modeling component then modifies the appropriate model parameters in an offline process to bring the operational filter back into optimal operation, and the autonomous monitor is reset to process new data. Chaer and Bishop⁴ proposed a genetic algorithm filter adaptation that used the HME as the performance function for parameter modification. However, this approach was found to evaluate performance relative only to filters in the HME, and Ely et al.⁵ reformulated the adaptive component to use an absolute performance index based on a sample statistic of filter residual histories.

The remainder of this paper is organized by a discussion of the architecture and theory of the HME, a discussion of experimental setup, experimental results, and a brief conclusion. Experiments were performed to test the ability of the autonomous monitor to detect and characterize 1) small unmodeled impulsive velocity changes, 2) mismodeling in the SRP model and unmodeled impulses, and 3) unmodeled impulses and measurement noise increases occurring in the same data span. All filters used in this work are extended Kalman filters (EKF)s^{3,6} processing MPF two-way Doppler Deep Space Network (DSN) data. The dynamic and measurement models used to define each filter in the HME are variations of the operational model developed by the MPF navigation team and detailed in the postmission navigation report.¹

Autonomous Monitoring Architecture

In the general HME configuration (Fig. 2 illustrates the specific configuration used here), each filter contains a unique measurement and dynamic model that can be represented by the parameter

Presented as Paper 2000-3939 at the AIAA/AAS Astrodynamics Specialist Conference, Denver, CO, 14–17 August 2000; received 6 February 2001; revision received 23 July 2001; accepted for publication 26 July 2001. Copyright © 2001 by the authors. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/02 \$10.00 in correspondence with the CCC.

*Aerospace Engineer, Aeroscience and Flight Mechanics Division. Member AIAA.

†Professor, Center for Space Research. Associate Fellow AIAA.

‡Staff Engineer, Navigation and Mission Design Section. Senior Member AIAA.

vector α . Thus, a particular filter realization may be thought of as an expert navigation system for a specific region of the modeling parameter space. The HME model is composed of macromode characterization banks (groups of filters) on the top level, which are in turn composed of micromode characterization experts (filters). Macromode environment changes considered here are unmodeled impulsive velocity changes, SRP model changes, and measurement system noise changes. Filter realizations representing macromode changes in the spacecraft environment are collected into competing

banks on the top level. The HME is separated into four filter banks. Three of the banks represent macromode changes in the spacecraft environment and contain multiple filters that model specific impulsive events, SRP model changes, and noise level changes. The fourth bank contains only the operational filter as an experimental control and should be selected as the most appropriate macromode by the HME in the absence of unmodeled events in the spacecraft environment.

The HME is regulated by two levels of single-layer neural networks known as gating networks (GNs) as indicated in Fig. 2. The bank-level GNs assign weights to filters within a single bank to indicate the relative residual performance of each realization. The top-level GN assigns weights to each bank according to the collective residual performance of the filters within the bank. Optimal filtering is assumed as long as the maximum value of the top-level gating vector \mathbf{g}_k at the time t_k of a measurement z_k is associated with the control bank, that is, $g_{3,k}$ is the maximum element of \mathbf{g}_k . A shift to suboptimal filtering is indicated when the maximum element of \mathbf{g}_k is associated with one of the banks representing an environment change macromode. Such a shift could be the result of a discrete event, as in the case of a thruster misfire, or a continuous or sustained environment change, as in the case of a steady change in shading of surfaces exposed to SRP. From a practical standpoint, a thruster misfire might only require a few hours worth of tracking data to manifest a recognizable degradation in operational filtering performance, whereas a subtle change in SRP effects might take several days or weeks of measurement data to present a degraded tracking solution. Characterization is accomplished by inspection of the macromode consistently given the highest top-level gating weight. Once the environment change macromode has been characterized, the adaptive modeling component then operates in an offline mode to adjust the parameter set by a genetic algorithm. Finally, the operational filter is updated and the autonomous monitoring process resumes with the adapted optimal operational filter.⁵

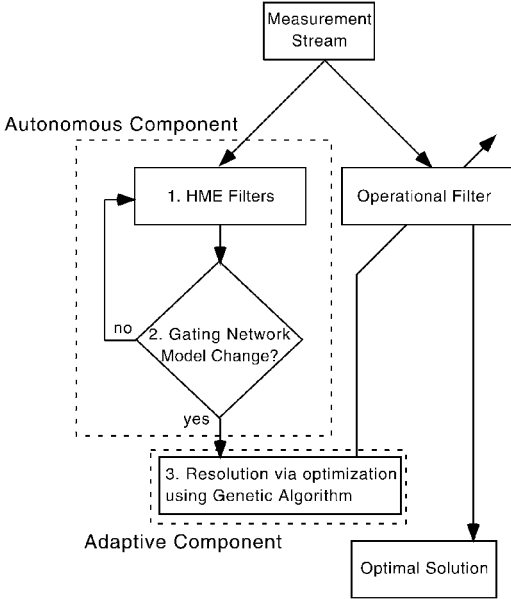


Fig. 1 Autonomous/adaptive navigation architecture.

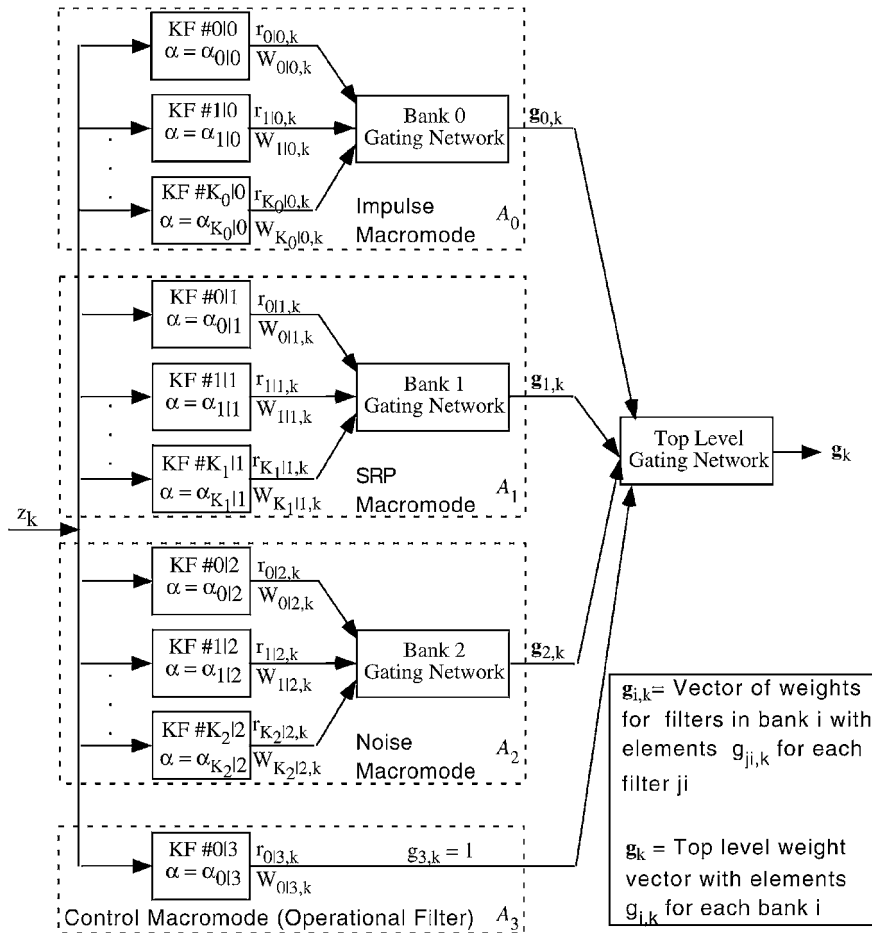


Fig. 2 HME configuration.

HME

The operation of the HME model regulated by multilevel GNs is based on maximizing the probability density of the current input measurement z_k as approximated by a weighted combination of the conditional probabilities of each of the L macromodes in the HME

$$f(z_k | Z_{k-1}) = \sum_{i=1}^L f(z_k | A_i, Z_{k-1}) g_{i,k}, \quad g_{i,k} = P(A_i | Z_{k-1})$$

where A_i is the set of all filter realizations in the i th macromode,

$$A_i = \{\alpha_{ji} | 0 < j \leq K_i, 0 < i \leq L\}$$

K_i is the number of filters in the i th bank, and Z_{k-1} represents the set of all earlier measurements. The conditional probability of the i th macromode is $f(z_k | A_i, Z_{k-1})$ and $P(A_i | Z_{k-1})$ is the a priori probability that the i th macromode models the input generating environment. The top level GN approximates the a priori probability using the gating weight $g_{i,k}$.

The conditional probability of the i th macromode is given by

$$f(z_k | A_i, Z_{k-1}) = \sum_{j=1}^{K_i} f(z_k | \alpha_{ji}, Z_{k-1}) g_{ji,k}$$

$$g_{ji,k} = P(\alpha_{ji} | Z_{k-1})$$

where $f(z_k | \alpha_{ji}, Z_{k-1})$ is the conditional probability of the j th filter (defined by α_{ji}) and $P(\alpha_{ji} | Z_{k-1})$ is the a priori probability among filters in the i th bank that the j th filter models the input generating environment. The bank-level GN approximates this a priori probability in the gating weight $g_{ji,k}$. If a Gaussian distribution is assumed on the measurement residuals of each filter, the probability distribution of the j th filter is given by

$$f(z_k | \alpha_{ji}) = \frac{1}{\sqrt{2\pi} W_{ji,k}} \exp\left(-\frac{r_{ji,k}^2}{2W_{ji,k}}\right)$$

where $r_{ji,k}$ is the prefit measurement residual and $W_{ji,k}$ is the innovations covariance of the j th filter at time t_k . The prefit measurement residual is given by

$$r_{ji,k} = z_k - h(\hat{x}_{ji,k}^{(-)}; \alpha_{ji})$$

where $h(\hat{x}_{ji,k}^{(-)}; \alpha_{ji})$ is the measurement computed from the Kalman filter model evaluated at the current state estimate $\hat{x}_{ji,k}^{(-)}$. The innovations covariance is a function of the state error covariance $P_{ji,k}^{(-)}$, the measurement mapping matrix $H_{ji,k}$, and the measurement noise covariance $\sigma_{ji,k}^2$,

$$W_{ji,k} = H_{ji,k} P_{ji,k}^{(-)} H_{ji,k}^T + \sigma_{ji,k}^2$$

The prefit residual, the innovations covariance, and the measurement noise covariance are all scalar values in this implementation because the filter is processing scalar Doppler observables.

The likelihood function of the hierarchy may, therefore, be expressed as a function of the individual filter conditional probabilities and the top- and bank-level gating weights (a priori probabilities):

$$f(z_k) = \sum_{i=1}^L g_i \sum_{j=1}^{K_i} f(z_k | \alpha_{ji}) g_{ji,k} \quad (1)$$

Note that explicit dependence on earlier information Z_{k-1} , will now be implicitly assumed to simplify the notation.

The gating weight a priori approximations are calculated by transforming the synaptic weights of the top- and bank-level GN neurons through the softmax operator.^{7,8} The gating weight of the i th output

Table 1 Softmax illustration

Bank	a_i	g_i
0	-2.0	0.003
1	0.0	0.0350
2	2.0	0.2583
3	3.0	0.7020
Sum	3.0	1.0

cell of the top-level GN is defined using softmax as a function of the top-level GN synaptic weights

$$g_i = \frac{e^{a_i}}{\sum_{j=1}^L e^{a_j}}$$

where a_i is the top-level synaptic weight for each bank-associated neuron.^{9,10} The gating weight assigned to the j th filter within the i th filter bank is similarly given by

$$g_{ji,k} = \frac{e^{a_{ji,k}}}{\sum_{n=1}^{K_i} e^{a_{ni,k}}}$$

The softmax function serves to provide a differentiable activation function that preserves rank order and generalizes a winner takes all paradigm by exponentially separating gating weights. As an example, consider the synaptic and gating weights in Table 1. Rank order is preserved because the synaptic weights of the banks map into the same magnitude order in the gating weights. The winner takes all paradigm is observed when the relative gating weight magnitudes within the rank order are shifted to favor the bank with the largest synaptic weight. For example, the 50% increase in synaptic weight from bank 2 to bank 3 translates into a marked 271% increase in gating weight.

The gating weights are renormalized with each adaptation of the synaptic weights, which occurs when a measurement is processed by the Kalman filters, and have the following properties:

$$0 < g_i < 1, \quad 0 < g_{ji} < 1 \quad (2)$$

$$\sum_{i=1}^L g_i = 1, \quad \sum_{j=1}^{K_i} g_{ji} = 1 \quad (3)$$

The gating weights may, therefore, be reasonably interpreted as the approximation to the conditional a priori probabilities¹¹ that each filter or bank of filters, depending on the level in the hierarchy, models the spacecraft dynamics and measurement environment correctly.

Adapting the synaptic weights (and, hence, modifying the gating weights) of the HME to maximize Eq. (1) is equivalent to adapting the synaptic weights to maximize the log-likelihood function

$$l = \ln f(z_k) \quad (4)$$

which proves easier to manipulate when evaluating a gradient with respect to the synaptic weights.

Substitution of the g into Eq. (4) gives the log likelihood in terms of the filter products and the gating weight vectors. The sensitivity of the log likelihood to the synaptic weights is

$$\frac{\partial l}{\partial a_{i,k}} = h_{i,k} - g_{i,k} \quad (5)$$

at the top level and

$$\frac{\partial l}{\partial a_{ji,k}} = h_{i,k} (h_{ji,k} - g_{ji,k}) \quad (6)$$

at the bank level. The h are defined as

$$h_{i,k} = \frac{g_{i,k} \sum_{j=1}^{K_i} f(z_k | \alpha_{ji}) g_{ji,k}}{\sum_{n=1}^L g_{n,k} \sum_{j=1}^{K_n} f(z_k | \alpha_{jn}) g_{jn,k}}$$

for the top level and as

$$h_{ji,k} = \frac{f(z_k | \alpha_{ji}) g_{ji,k}}{\sum_{j=1}^{K_i} f(z_k | \alpha_{ji}) g_{ji,k}}$$

for the j th filter on the bank level. Because of their dependence on filter products, the h may be thought of as a posteriori probabilities of each filter or bank. It is through the a posteriori probabilities that the filter performances influence the GN learning process. This a posteriori probabilistic interpretation of filter outputs has been used in other adaptive Kalman filter applications with varying degrees of success.^{12,13}

The GN updates the synaptic weights $a_{i,k}$ and $a_{ji,k}$ after the experts process z_k by the gradient ascent procedure, which seeks to maximize Eq. (4) (Ref. 9). The adaptation of the synaptic weights is

$$a_{i,k+1} = a_{i,k} + \eta \frac{\partial l}{\partial a_i} = a_{i,k} + \eta(h_{i,k} - g_{i,k}) \quad (7)$$

$$a_{ji,k+1} = a_{ji,k} + \eta \frac{\partial l}{\partial a_{ji,k}} = a_{ji,k} + \eta h_{i,k}(h_{ji,k} - g_{ji,k}) \quad (8)$$

where η is a learning rate parameter that is not generally the same for the top-level or individual bank-level GNs. The effective result of this learning scheme is that the gating weights track the a posteriori probabilities, which are functions of filter residual and innovations covariance magnitudes. When the synaptic weights are adapted to maximize Eq. (4), the best performing filters will generally be given a weight close to unity by the bank-level GN. Similarly, the best performing bank will be given a weight close to unity by the top-level GN; however, the presence of sufficiently poor performing filters in a bank containing the optimal filter may cause that bank to not receive the highest top-level gating weight. Note also that the bank-level learning rule in Eq. (9) scales the update to the synaptic weight by the h value for that bank. Because of this scaling, internal learning progresses at a slower rate for filter banks that are assigned low top-level gating weights.

Experimental Setup

The experimental application of the autonomous monitoring component is the MPF cruise segment beginning after trajectory correction maneuver 2 on 4 February 1997 and ending 17 April 1997. The two-way DSN Doppler data processed from this period is collected every 10 min in roughly 8-h passes from each DSN station. Nongravitational accelerations, SRP, and measurement noise have been identified as three of the most significant sources of uncertainty during MPF cruise.¹⁴ Chaer et al.² were able to characterize successfully changes in process noise statistics related to these error sources using simulated MPF data and a linear Kalman filter implemented with an approximate dynamic model. Their results indicated that an HME implementation using high-fidelity filters processing real data might be able to characterize deterministic changes in these error sources. Therefore, the goal of the following experiments is to determine if a general HME configuration can be used to characterize 1) nongravitational accelerations in the form of unmodeled impulsive maneuvers, 2) SRP environment changes, and 3) unmodeled impulses and measurement noise statistics. The top-level GN must clearly assign weight to the correct macromode characterization bank in the event of a spacecraft environment change. Also, the HME must be able to avoid false characterization by giving the control filter bank the highest weight in the absence of environment changes. An initial learning period is expected because all filters and banks are initially weighted equally in the HME. The length of this period is inversely related to the magnitude of the learning rates in the HME.

All filters in the HME are configured with the parameter realization defined in Ref. 1 unless otherwise noted. The initial covariances used in all filters are summarized in Table 2. The optimal SRP model is indicated as "Ely" and is a model tuned by the genetic algorithm in the adaptive modeling component of navigation architecture.⁵ The learning rate for all GNs is set to 10 in the following experiments. This value was chosen empirically to display both responsiveness

Table 2 Initial covariances for all filters

States	Dimension	1 - σ	Units
Position	3	1e4	km
Velocity	3	5	m/s
4 Feb. impulse	3	1e-1	cm/s
25 March/test impulse	3	1e-1	cm/s
Measurement bias	1	0.05	Hz
Dry troposphere	3	3.4e-10	—
Wet troposphere	3	1.65e-10	—
Area scale factors	6	1.0	—

Table 3 Impulsive event experiment HME configuration

Filter number	Impulse	SRP model	R
0, 0	4 Feb.	Ely	0.01
1, 0	22 Feb.	Ely	0.01
2, 0	12 March	Ely	0.01
3, 0	30 March	Ely	0.01
0, 1	—	MPF 2	0.01
1, 1	—	MPF 4	0.01
0, 2	—	Ely	0.003
1, 2	—	Ely	0.03
2, 2	—	Ely	0.09
0, 3	—	Ely	0.01

to environment changes and fidelity in the presence of noise. In practice, the gating calculations are much less demanding than the individual filter calculations, and a range of learning rates could be used on the same residuals. Therefore, the results that follow represent one possible sensitivity of the HME that might be available to navigators employing the autonomous monitor. The implications of tuning the learning parameter are beyond the scope of this work and are not presented here.

Results

Impulsive Events

To characterize unmodeled impulses, the HME must contain a specific bank of filters that will accommodate an unmodeled velocity change better than the operational filter in the control bank. Bank 0 in the HME configuration of Table 3 contains filters with impulsive maneuver filter states spaced evenly in time for the Doppler data span under consideration. These impulses are modeled as nominally zero magnitude, but the filter estimates corrections to the impulses applied at these test times. Therefore, additional state space is introduced to accommodate unmodeled velocity changes by absorbing Doppler residual signals through the Kalman gain for the test impulse estimate. The small MPF trajectory maneuver, approximately 0.7 mm/s, on 25 March 1997 will be used as the impulse to be detected and characterized in this experiment. This maneuver magnitude is just above the 0.18-mm/s noise level modeled in the operational filter Doppler measurements¹ and has been omitted from the model of all filters in the hierarchy.

The filters in bank 1 use the MPF navigation team SRP solutions 2 and 4 as detectors for SRP change and the operational Doppler noise value of 0.01 Hz (0.18 mm/s) (Ref. 1). Filters in bank 2 model changes in the Doppler measurement noise but are otherwise optimal. Filter (0, 3) is the operational filter and is optimal other than the absence of the 25 March impulse.

The top-level gating weights from this experiment are provided in Fig. 3, where the times of the unmodeled impulse and the bank 0 test impulse states are indicated by vertical lines. Note that the weights are assigned only at times of data availability and, therefore, have gaps between data passes. The gating weights are equally distributed between the impulse and control banks for the first 10 days of data. This result is to be expected because three of the filters in bank 0 contain the same dynamic model as the operational filter over the first quarter of the data span. However, when the 22 February test impulse in filter (1, 0) comes and goes without an impulse occurring, weight is reassigned from the impulse test bank to the control bank.

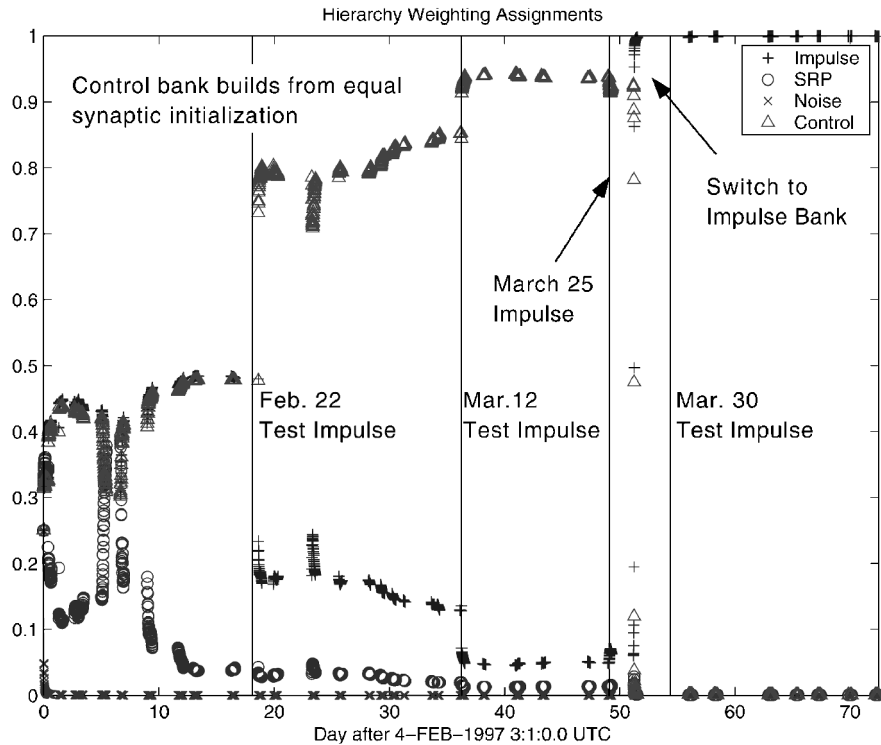


Fig. 3 Impulse event top-level weights.

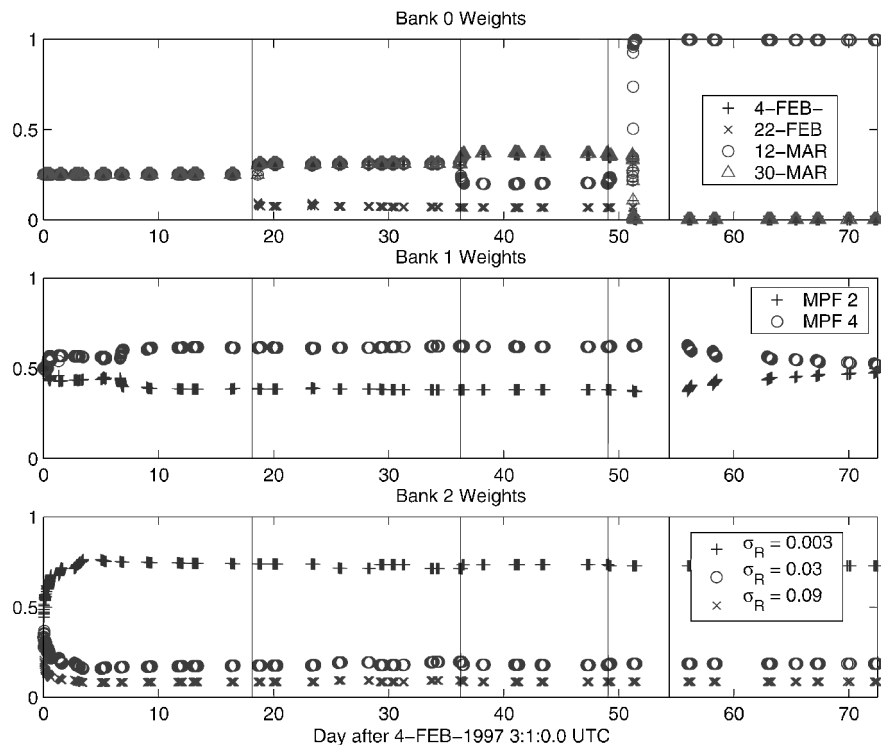


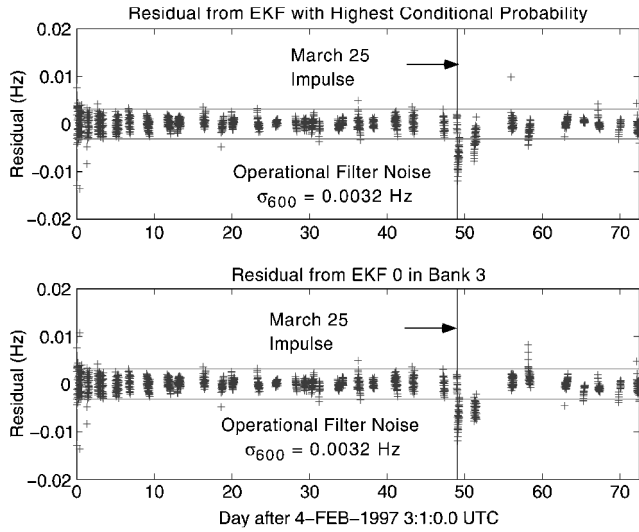
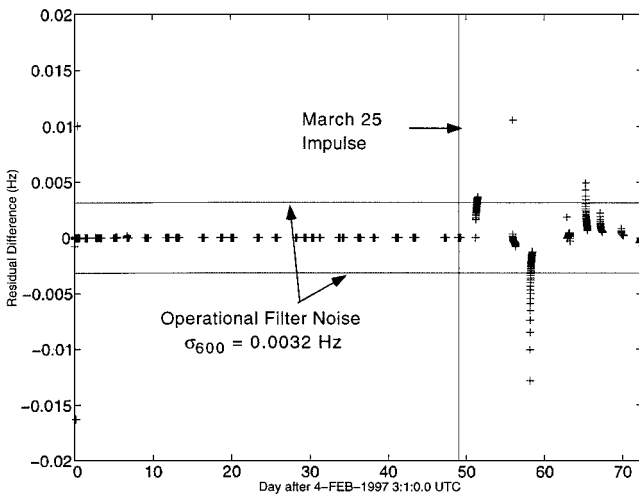
Fig. 4 Impulse event bank-level weights.

The autonomous monitor passes an important test at this point by avoiding a false detection of a macromode change. A similar effect is observed on 12 March to a smaller extent. The unmodeled impulse occurs in the middle of the data pass on day 49 but does not manifest a significant enough change in the data for characterization until the next pass when the impulse bank is correctly assigned the majority of the gating weight. A navigator examining this gating history would be justified in adapting the operational filter to accommodate an impulse between 12 March and 27 March. The bank-level gating weights in Fig. 4 reflect a similar behavior because filter (2, 0) is assigned the most weight in bank 0 on the pass after the unmod-

eled impulse. The SRP alternative models in bank 1 are weighted relatively the same within their bank. The small noise filter (0, 2) is selected throughout the data span within bank 2. However, recall from Eq. (9) that bank-level learning is scaled by the top level a posteriori weight $h_{i,k}$ associated with each bank. Because the gradient learning rule in Eq. (8) indicates that the gating weights track the a posteriori weights and the noise bank gating weight goes to zero in Fig. 3, it can be concluded that relatively no internal learning occurs in bank 2 because of its poor collective top-level performance. To illustrate the sensitivity of the HME selected filter and the

Table 4 SRP model change experiment HME configuration

Filter number	Impulse	SRP model	R
0, 0	4 Feb.	MPF 4	0.01
1, 0	22 Feb.	MPF 4	0.01
2, 0	12 March	MPF 4	0.01
3, 0	30 March	MPF 4	0.01
0, 1	—	MPF 2	0.01
1, 1	—	Ely	0.01
0, 2	—	MPF 4	0.003
1, 2	—	MPF 4	0.03
2, 2	—	MPF 4	0.09
0, 3	—	MPF 4	0.01

**Fig. 5** Impulse event operational and HME filter residuals.**Fig. 6** Impulse event operational and HME filter residual differences.

operational filter are presented in Fig. 5. The HME selected filter is the filter with highest bank-level weight in the bank with the highest top-level gating weight. Until the time of the unmodeled impulse, the HME filter produces residuals essentially equivalent to the control filter. Differences in the residuals do not begin to appear until day 51 (Fig. 6) and are not outside the measurement noise range until after the impulse bank has been selected on the top level. This indicates that the autonomous monitor characterized the unmodeled impulse when the residual signature was still within the noise 1σ .

SRP Model Changes

To investigate the ability of the autonomous component to characterize a continuous change in the SRP model, the HME was config-

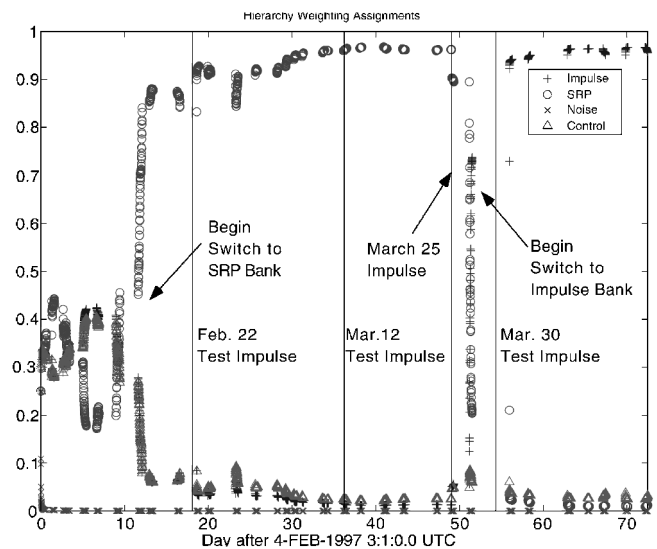
ured as detailed in Table 4. The basic structure is the same as in the impulsive event detection with the exception that all of the filters not involved with SRP change characterization utilize the MPF navigation team SRP model 4. The assumption of this scenario is that MPF SRP model 4 is incorrectly believed to be the optimal model and is therefore used in the operational filter and most other filters. The two filters in the SRP change bank 1 use the Ely genetic algorithm tuned SRP model and the MPF navigation team model 2 to provide alternative SRP dynamics to the operational model. The Doppler noise values are configured in the same manner as before with bank 2 filters modeling changes in the measurement noise. To examine the degree to which the top-level GN can switch between macromodes, the small 25 March impulse is again omitted from all filters.

After the initial 10-day learning period, the top-level gating weights are definitively assigned to the SRP model change bank as illustrated in Fig. 7. When the bank level weights in Fig. 8 are examined, it can be seen that the optimal Ely SRP model in expert (0, 1) receives the majority of weight from the beginning of the data span, but also gains a definitive weighting majority after day 10. The top-level SRP bank maintains the maximum weighting through the test impulses on 12 February and 22 March, but the impulse bank is selected beginning on the data pass after the unmodeled impulse. Therefore, in this experiment the autonomous monitor is able to distinguish between a continuous and a discrete environment change within the same data set.

The HME selected filter and operational filter residuals in Fig. 9 are nearly identical to visual inspection. Furthermore, the HME and operational filter residual differences in Fig. 10 are well below the magnitude of the nominal MPF Doppler noise. Therefore, the selection of the SRP macromode on the top level is based on the continuous performance of the filters in bank 1. This result is markedly different from the HME decisions involving unmodeled impulses in the previous experiment, where the residual difference between

Table 5 Impulse and noise experiment HME configuration

Filter number	Impulse	SRP model	R
0, 0	4 Feb.	Ely	0.01
1, 0	22 Feb.	Ely	0.01
2, 0	12 March	Ely	0.01
3, 0	30 March	Ely	0.01
0, 1	—	MPF 2	0.01
1, 1	—	MPF 4	0.01
0, 2	—	Ely	0.003
1, 2	—	Ely	0.03
2, 2	—	Ely	0.09
0, 3	—	Ely	0.01

**Fig. 7** SRP change top-level weights.

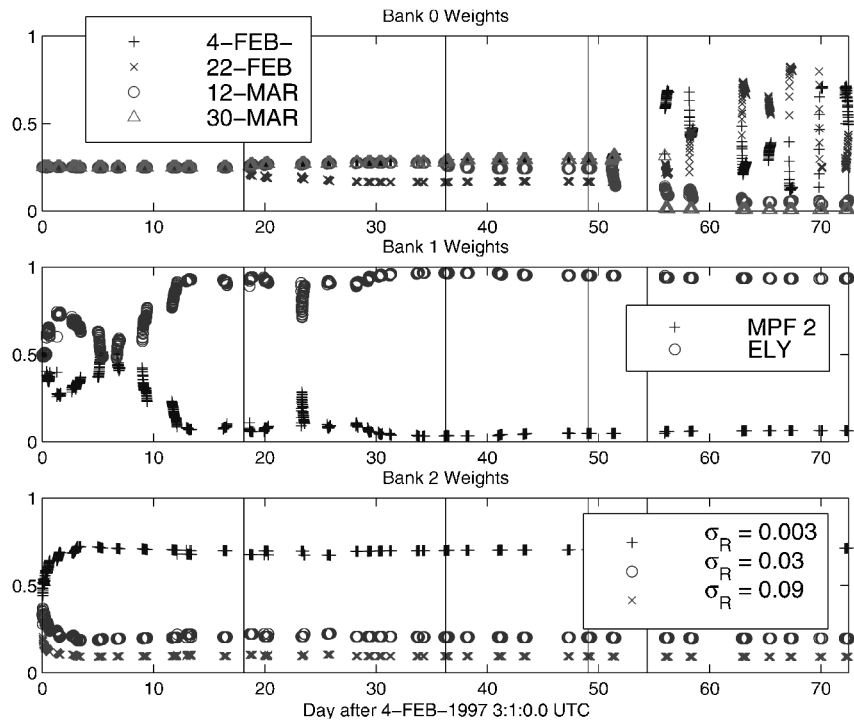


Fig. 8 SRP change bank-level weights.

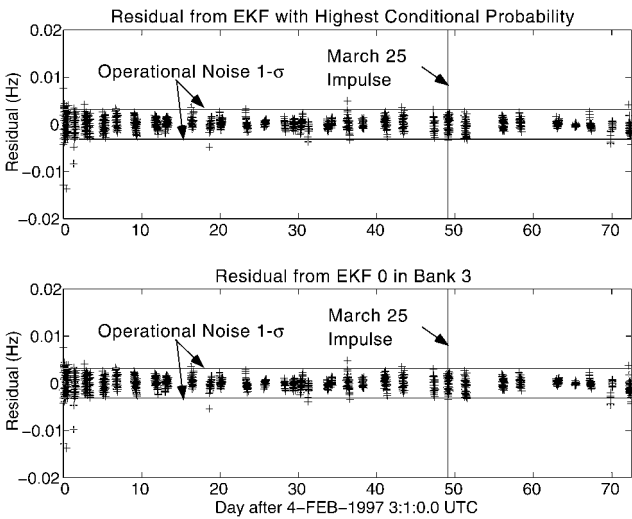


Fig. 9 SRP change operational and HME filter residuals.

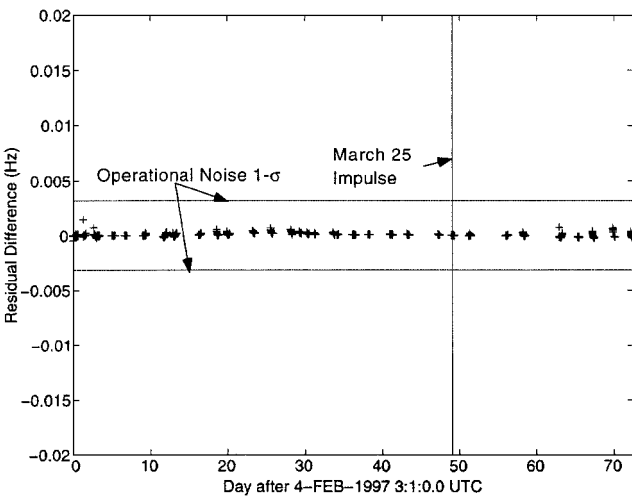


Fig. 10 SRP change operational and HME filter residual differences.

the operational and HME filters would demonstrate a rapid and significant magnitude change within a single pass. Because this experiment also includes the unmodeled impulse on 25 March, note that the HME correctly characterizes this event without a noticeable residual signature after the impulse on day 49 in Fig. 10.

Noise Change with Impulsive Event

The impulsive event detection performed in the first experiment is now conducted with the addition of the Gaussian noise signal plotted in Fig. 11. The addition of this noise from day 5 to day 15 gives the measurements an equivalent total noise value of 0.04 Hz for this time period. The HME is configured exactly the same as in the impulse detection experiment as illustrated in Table 5.

The top-level and bank-level weights for the noise and impulse characterization experiment are provided in Figs. 12 and 13. The top-level weights exhibit a brief initial learning period before the noise bank moves to almost unity, weighting very close to the beginning of the increased noise segment. On the bank level, filter (1, 2) with a noise statistic of 0.03 Hz moves to unity weighting

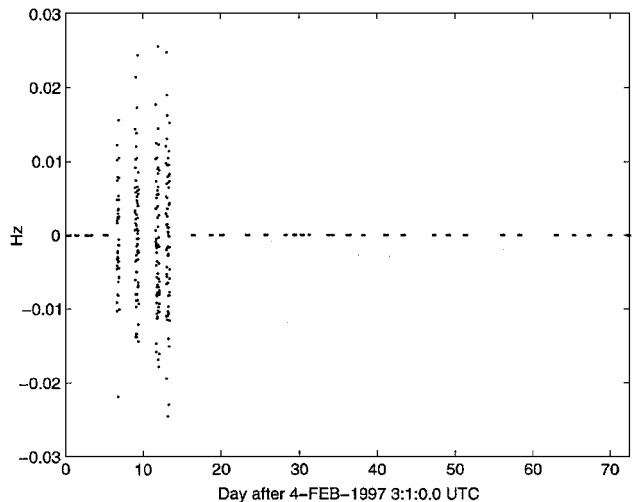


Fig. 11 Noise added to MPF Doppler data.

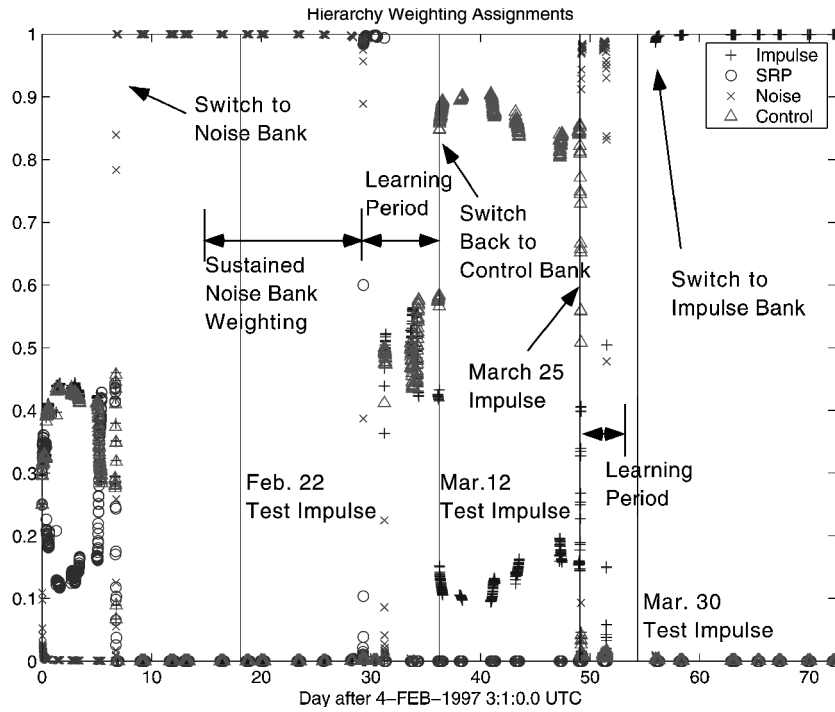


Fig. 12 Noise and impulse top-level weights.

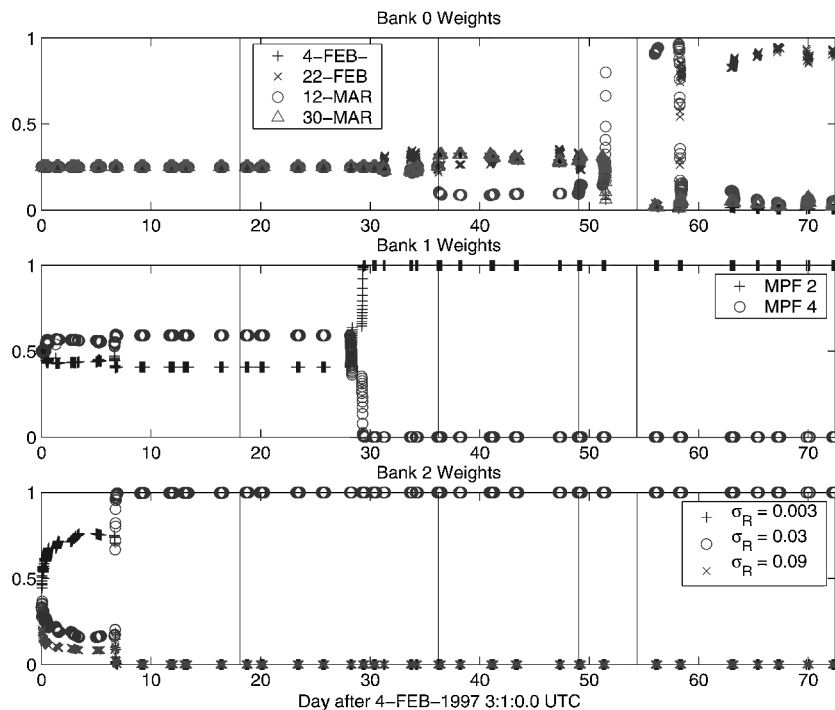


Fig. 13 Noise and impulse bank-level weights.

over the 0.003- and 0.09-Hz filters. The increased noise segment ends on day 15, but the noise bank maintains unity top-level weight until the learning period starting near day 30. It is not clear whether the extended weighting of the noise bank is a property of the GN or the EKF's or both. However, during the learning period, the control and impulse banks are weighted equally. Once the 12 March test impulse in filter (2, 0) passes near day 35, the control bank is definitively selected as the best performer. When the unmodeled impulse occurs on 25 March, a second learning period occurs where the noise bank is given the highest weight for two data passes. After 30 March, the impulsive filter is given unity weighting. The soft decision capa-

bility of the multilevel HME is demonstrated in this switch because the 12 March test impulse filter is selected after the unmodeled impulse, only to give way to the 30 March test impulse filter on the bank level, but the top-level decision is nonetheless correct.

The residual comparisons between the HME selected filter and the operational filter presented in Figs. 14 and 15 show that there is only a small difference in the residuals from day 25 to day 30, indicating that the noise bank filters are performing as well as the operational filter during this span. After the unmodeled impulse on 25 March, there is a two-pass transition before the top-level weights are assigned to the impulse detection bank.

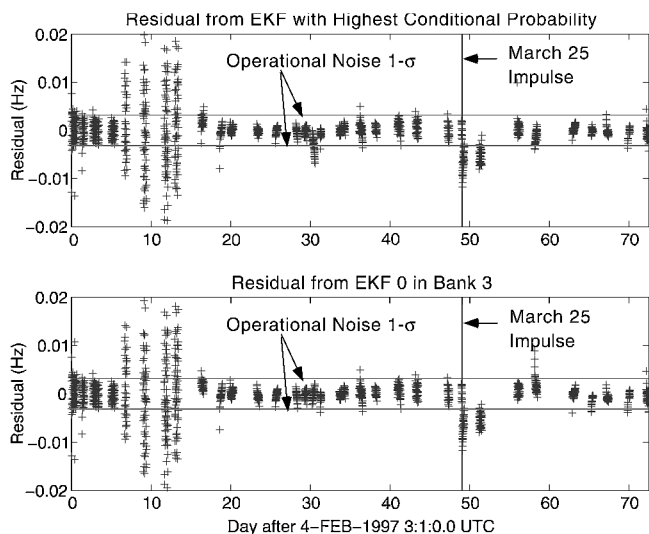


Fig. 14 Noise and impulse operational and HME filter residuals.

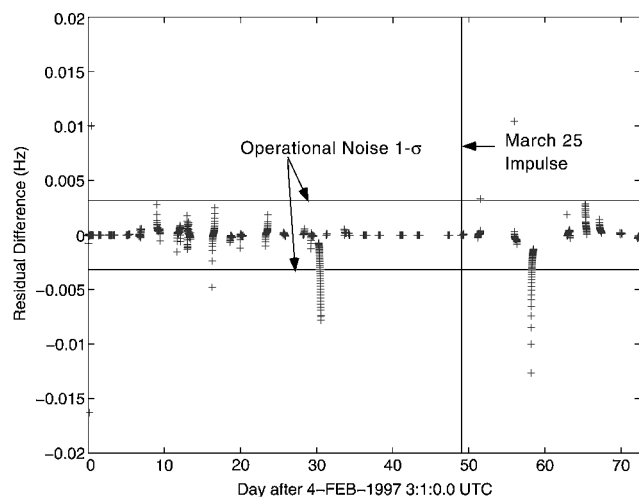


Fig. 15 Noise and impulse operational and HME filter residual differences.

Conclusions

The autonomous monitoring algorithm successfully detected and characterized unmodeled macromode environment changes in the MPF interplanetary cruise environment within one or two data passes for single macromode unmodeled environment changes. Unmodeled events were characterized correctly in all three experiments, and false characterizations in the absence of environment changes were not significantly present. Learning periods were observed at the beginning of each experiment when no filter bank was clearly selected or the incorrect filter bank was temporarily selected by the top-level GN. The multimode MPF environment changes in the third experiment created two distinct learning periods within the data span after each mode change and incorrectly maintained noise macromode selection for approximately 15 days after the increased noise period elapsed. Some learning is required for the top-level GN to make robust decisions, otherwise a single measurement could cause an instantaneous reassignment of gating weight and lead to false detection or incorrect characterization of environment changes. The unmodeled impulsive maneuver and solar radiation pressure mismodeling were both characterized when the difference between the operational and HME selected measurement residuals were still below the level of the Doppler DSN measurement noise.

The learning periods at the beginning of the data span might be reduced if the GNs were initialized with synaptic weights favoring the operational filter in the control bank instead of equal initial weighting as in these experiments. Such a biased initial weighting scheme has intuitive merit because the operational filter is believed

to be the best model of the spacecraft environment when the first measurement is processed. However, because the adaptation of the synaptic weights is sensitive to their magnitude during application of the gradient learning rule, care must be taken to not overlook changes that occur shortly after the data span begins. As always, the balance between response time and robust decision making must be observed.

Further experimentation is required to determine the effectiveness of the impulse characterization configuration used in bank 0. In particular, the temporal spacing of test impulse filter states required to determine adequately the approximate time of unmodeled discrete velocity changes in an operational setting has yet to be determined. It is possible that the adaptation of the operational filter to account for unmodeled impulses might require the genetic algorithm to include event time as one of its parameters. Previous work with a single linear Kalman filter modeling multiple unknown accelerations inspires the creation of a queue of impulse detection EKFs in bank 0. In the queued approach, a manageable number of filters populate bank 0 with a reasonable temporal test impulse spacing. If bank 0 does not accrue top-level weight within a certain time interval after the first test impulse in the queue has occurred, then that filter is removed from the bank and a new filter with an upcoming test impulse replaces it at the end of the queue. The multilevel structure of the HME would be useful in this approach because bank-level probabilities would be reassigned equally when the queue changes, but the top-level weight assigned to the impulse detection bank would remain the same. This approach is appealing because it limits the number of active filters required for small temporal spacing of test impulses in bank 0.

Several operational questions must be addressed before the autonomous monitoring and the offline adaptive components can be fully integrated. Given that periods of learning can occur even with acceptable top-level decision making, two distinct problems arise. First, a method of distinguishing ongoing learning from environment changes beyond the scope of those modeled in the HME configuration of filters should be developed. It is possible that the HME can diverge if a change occurs in which no filter or bank correctly models the spacecraft environment. Large residuals in all filters can cause a numerical underflow problem in the a posteriori probability calculations. A time limit or data quantity limit on learning along with minimum residual magnitude monitoring should adequately address this issue. Second, a method of determining when adaptation should take place should be developed. The operational filter should be adapted by the genetic algorithm only after the top-level weights have maintained a threshold value for a specified number of data points. This would avoid adaptation in response to transient or learning period weighting regimes.

Finally, it is envisioned that the autonomous monitoring component is a candidate for implementation onboard a spacecraft during interplanetary cruise. The primary concern in such an implementation would be the computational load required for running multiple environment realizations in parallel. However, the Kalman filter is a sequential estimation algorithm known for its efficiency and should easily be accommodated by modern flight computers during nominal cruise operations. The offline adaptation component is more demanding of computer resources and may have to be relegated to ground-based operation for difficult parameter optimizations. A secondary modification for onboard operation would be processing data terminating at the spacecraft rather than two-way data types that originate and terminate at a ground station. Such data could be two-way data relayed from spacecraft to ground station to spacecraft or one-way data transmitted from the ground station. Because the HME formulation only requires the statistical evaluation provided by residual analysis, other data types such as optical data originating from sensors onboard the vehicle are theoretically possible for use in the autonomous monitoring scheme.

Acknowledgments

The research described in this paper was performed in part by the Jet Propulsion Laboratory, California Institute of Technology, under a Contract with NASA, and was supported in part by a Grant from the Center for Space Research at the University of Texas at Austin.

References

¹Vaughan, R., Kallemeyn, P., Spencer, D., and Braun, R., "Navigation Flight Operations for Mars Pathfinder," American Astronautical Society, AAS Paper 98-145, Feb. 1998.

²Chaer, W., Bishop, R., and Ghosh, J., "Hierarchical Adaptive Kalman Filtering for Interplanetary Orbit Determination," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 34, No. 3, 1998, pp. 1-14.

³Crain, T., and Bishop, R., "The Mixture-of-Experts Gating Network: Integration into the ARTSN Extended Kalman Filter," Center for Space Research, TM CSR-TM-99-01, Univ. of Texas at Austin, Austin, TX, March 1999.

⁴Chaer, W., and Bishop, R., "Adaptive Kalman Filtering with Genetic Algorithms," *Advances in the Astronautical Sciences*, edited by R. Proulx, J. Liu, P. Siedelmann, and S. Alfano, Vol. 89, Univelt, San Diego, CA, Pt. 1, 1995, pp. 141-156.

⁵Ely, T., Bishop, R., and Crain, T., "Adaptive Interplanetary Navigation Using Genetic Algorithms," *Journal of Astronautical Sciences*, Vol. 48, Nos. 2-3, 2000, pp. 287-303.

⁶Gelb, A., *Applied Optimal Estimation*, MIT Press, Cambridge, MA, 1974, pp. 102-155, 180-228.

⁷Bridle, J., "Probabilistic Interpretation of Feedforward Classification Network Outputs, with Relationship to Special Pattern Recognition," *Neurocomputing: Algorithms, Architectures, and Applications*, edited by

F. Fouglerman-Soulie and J. Hérault, Springer-Verlag, New York, 1990, pp. 227-236.

⁸Bridle, J., "Training Stochastic Model Recognition Algorithms as Networks can Lead to Maximum Mutual Information Estimation of Parameters," *Advances in Neural Information Processing Systems 2*, edited by D. Touretzky, Morgan Kaufmann, San Mateo, CA, 1990, pp. 211-217.

⁹Haykin, S., *Neural Networks a Comprehensive Foundation*, 2nd ed., Prentice-Hall, Upper Saddle River, NJ, 1999, pp. 368-373.

¹⁰Crain, T., and Bishop, R., "Unmodeled Impulse Detection and Identification During Mars Pathfinder Cruise," Center for Space Research, TM CSR-TM-00-01, Univ. of Texas at Austin, Austin, TX, March 2000.

¹¹Haykin, S., *Neural Networks a Comprehensive Foundation*, 1st ed., Macmillan, New York, 1994, pp. 473-494.

¹²Gholson, N., and Moose, R., "Maneuvering Target Tracking Using Adaptive State Estimation," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 13, No. 3, 1997, pp. 310-317.

¹³Burkhart, P., and Bishop, R., "Adaptive Orbit Determination for Interplanetary Spacecraft," *Journal of Guidance, Control, and Dynamics*, Vol. 19, No. 3, 1997, pp. 693-701.

¹⁴Burkhart, P., Bishop, R., and Estefan, J., "Covariance Analysis of Mars Pathfinder Interplanetary Cruise," *Advances in the Astronautical Sciences*, edited by R. Proulx, J. Liu, P. Siedelmann, and S. Alfano, Vol. 89, Univelt, San Diego, CA, Pt. 1, 1995, pp. 267-285.