

Gyro-based maximum-likelihood thruster fault detection and identification¹

Edward Wilson², Chris Lages², and Robert Mah³

Smart Systems Research Laboratory, NASA Ames Research Center, MS 269-1, Moffett Field, CA 94035

Abstract

When building smaller, less expensive spacecraft, there is a need for intelligent fault tolerance vs. increased hardware redundancy. If fault tolerance can be achieved using existing navigation sensors, cost and vehicle complexity can be reduced. A maximum-likelihood-based approach to thruster fault detection and identification (FDI) for spacecraft is developed here and applied in simulation to the X-38 space vehicle. The system uses only gyro signals to detect and identify hard, abrupt, single- and multiple-jet on- and off-failures. Faults are detected within one second and identified within one to five seconds.

1. Introduction

The FDI system presented here was developed through application to two specific thruster-controlled spacecraft presently under development at NASA Johnson Space Center: the X-38 [12] and the Mini-AERCam. Its application to the X-38, shown in Figure 1, is presented in this paper.




Figure 1: x38, with entry vehicle and de-orbit propulsion stage [7]

The Crew Return Vehicle (CRV) consists of a manned space vehicle, the Entry Vehicle (EV), based on a lifting-body design, and a De-orbit Propulsion Stage (DPS). The

CRV is designed to remain docked to the space station in a dormant mode for several years until needed by the crew in an emergency. The X-38 (vehicle 201) is the unmanned test vehicle for the CRV. Both vehicles are designed to maneuver on-orbit, de-orbit, and land using a large parafoil.

The DPS includes a set of axial and reaction control system (RCS) thrusters fed by three mono-propellant hydrazine tanks. Although the CRV will have pressure sensors in the thrusters to detect failures, the X-38 has only temperature sensors. In this research, a fault detection and identification (FDI) system is developed that uses only gyro signals (angular rate measurements) to detect and identify abrupt, single- and multiple-jet, hard-on or hard-off thruster failures.

1.1 Related research

Several FDI approaches reported in the literature [4] perform well on a variety of applications. However, the on-off nature of the thrusters present in the class of applications addressed here limits the viability of many general-purpose methods. For example, if a thruster has failed off, it will appear to be working correctly at all times that it is not commanded to fire. This paper presents a general approach for this class of problems that has been validated through application to specific, realistic spacecraft applications.

Deyst and Deckert [2] developed a maximum-likelihood based approach for detecting leaking thrusters for the Space Shuttle orbiter's RCS jets. The method for detecting soft failures was also extended to detect hard RCS jet failures. The maximum-likelihood method presented in that work is used and extended in this research.

Wilson and Rock [10] [11] developed an FDI method based on exponentially weighted recursive least squares estimation using accelerometer and angular rate sensors. A neural network then provided adaptive control reconfiguration to multiple destabilizing hard and soft thruster failures. This was applied to a 3-degree-of-freedom air-bearing vehicle.

2. Problem definition

Hard, abrupt, thruster failures resulting from a single point of failure (in valves, plumbing, electronics, etc.) are monitored. These can include single- or (simultaneous)

¹ This research was funded by NASA Headquarters, HQ AA, PWC 349-00.

² Drs. Wilson and Lages are with Intellization, ed.wilson and chris.lages@intellization.com

³ Dr. Mah is with NASA Ames Research Center, rmah@mail.arc.nasa.gov

multiple-jet failures in either a failed-on or failed-off condition. The DPS has 8 axial thrusters (500 Newtons thrust level each) that fire along the longitudinal axis of the vehicle, providing the required de-orbit thrust for the 13,600 kg vehicle. During the 8-to-15-minute de-orbit burn, six of the eight thrusters fire continuously, controlled open loop, with the six chosen symmetrically to produce minimal torque on the vehicle. The DPS also contains 8 RCS thrusters (106 Newtons thrust level each) that are fired in sets of two or four by the attitude control system to control the roll, pitch, and yaw about the body axes. The EV has a completely separate set of RCS thrusters for use after DPS separation – those are not considered here. Figure 2 is a rear-view schematic of the X-38 showing EV RCS, DPS RCS, and DPS axial thrusters. The axial thrusters fire directly back along the x-axis, and the DPS RCS thrusters fire in the y-z plane, with no x-axis component.




Figure 2: X-38 thruster configuration

Accelerometers and ring-laser gyros in the Honeywell Space Integrated GPS / INS (SIGI) [9] are available for monitoring vehicle motions. Temperature sensors in the thrusters provide failure information as well, but the response time limits the ability to detect failures sufficiently quickly – they rise in about one second, but cool over a period of minutes. At this point, the Fault Detection and Identification (FDI) system has been developed without using temperature information, but it could be added at a later point. Thruster faults will be detected by comparing vehicle motions (at this point, only rotational accelerations are used, but translational accelerations would enable even more accurate FDI) to the vehicle motions that would result if certain failures have occurred.

2.1 Equations of motion

Starting with Euler’s dynamical equation, and assuming the spacecraft inertia matrix is constant, the rotational equations of motion (EOM) are [1]

$$\dot{\mathbf{w}} = I^{-1}(\mathbf{t} - \mathbf{w} \times I\mathbf{w})$$

where I is the spacecraft inertia matrix, \mathbf{w} is the angular velocity of the body-fixed frame with respect to an inertial reference frame, and \mathbf{t} is the sum of all torques on the body.

2.2 Simulation

Several random variations are added to this dynamic model, including (values given are the 3-sigma value of a Gaussian distribution about the true or nominal value): pulse-to-pulse thruster strength variability of 15%; constant thruster strength bias of 5%; inertia matrix elements constant bias of 5%; constant mass bias of 1%; and center of mass (CM) location offset of 5 mm along x and y-axes and 25 mm along the z-axis⁴. These values are all conservative estimates (i.e., at least as large as the actual) based on the actual X-38 design.

A dynamic simulation was developed using MATLAB [6]. As with the X-38 design, the control loop runs at 10 Hz, and unfiltered gyro data is read at 50 Hz. The FDI runs at 10 Hz. A controller that regulates to a commanded attitude calculates the thruster commands; the EOM from above, including the random variations, are integrated; the FDI system detects and identifies failures; and a MATLAB-based visualization displays the vehicle status and FDI results as shown in Figure 4.

3. Fault detection and identification

It is generally true in system identification (ID) or FDI systems that reducing the degrees of freedom to be considered or otherwise constraining the problem will improve identification or detection performance. As will be discussed in Section 3.10, some alternative approaches were initially used to solve this problem that attempted to identify the strengths of the un-failed thrusters as well as finding the failures. This approach worked well on a simplified version of the problem, but became unreliable when all 16 thrusters were present, both on and off failures were considered, mass properties were allowed to vary within tolerance, and in the presence of gyro noise. This led to the approach described below, which solves the problem taking full advantage of the problem statement – namely that only a single failure mode from a finite list of candidates can be present, and that it will appear abruptly.

3.1 Summary of the algorithm, nomenclature

At every control update, the disturbing acceleration, $\hat{\mathbf{a}}_{disturbing}$, is calculated. This vector is compared with the vector of disturbing angular accelerations corresponding to each possible failure mode. After a fault is detected, and once a clear match is found (the likelihood is sufficiently higher than all other possibilities), the failure mode is identified. Specifics regarding filtering and other calculations follow.

⁴ The SIGI gyro noise spec was used, but is not published here

active – describes individual failure modes at each control update. An “off” (“on”) failure is said to be *active* if the corresponding thruster is (is not) commanded to fire during that sample period.

D – [3-by- N matrix] unit vectors indicating the direction of thrust in the body frame

F_{nom} – [N-by-N diagonal matrix] nominal strength of each thruster at full tank pressure

$F_{nom,k}$ – [N-by-1 vector] nominal force from each thruster at time step k , accounting for estimated blowdown and firing commands

i – failure mode number

inactive – opposite of *active*

I_{nom} – [3-by-3 matrix] nominal spacecraft inertia tensor

k – control (and FDI) update counter

L – [3-by- N matrix] x - y - z location of each thruster in the body frame (changes with the center of mass location)

N – number of thrusters, 16 for the X-38

P_a – [3-by-3 matrix] estimation error covariance of the disturbing acceleration

$T_{command,k}$ – [N-by-1 vector of 1’s and 0’s] which thrusters are commanded to fire at time step k

$\hat{\mathbf{a}}$ – measured (estimated) vehicle angular acceleration

$\mathbf{a}_{nom-system}$ – nominal-system acceleration - the angular acceleration that should result if no failures are present and all physical parameters are at their nominal value

$\hat{\mathbf{a}}_{disturbing}$ – [3-by-1 vector] measured disturbing acceleration, $\hat{\mathbf{a}}_{disturbing} = \hat{\mathbf{a}} - \mathbf{a}_{nom-system}$

$\mathbf{a}_{disturbing,i}$ – disturbing acceleration vector corresponding to failure mode, i , based on nominal values. For example, failure mode #1 corresponds to RCS jet 1 being failed off, and $\mathbf{a}_{disturbing,1}$ is [-0.0144, -0.0015, 0.0045] rad/sec² for the body roll pitch and yaw axes. This means that if RCS jet 1 is commanded to fire, and it has failed off, $\hat{\mathbf{a}}_{disturbing}$ should equal $\mathbf{a}_{disturbing,1}$

$\mathbf{a}_{disturbing}^*$ – [3-by-1 vector] disturbing acceleration vector corresponding to the true failure mode

$\Gamma_{blowdown}$ – a scalar multiplier representing the reduction in thrust with reduced tank pressure

$\lambda_{active,i}$ – likelihood argument for failure mode i , based on times when the failure mode is *active*

$\lambda_{inactive,i}$ – likelihood argument for failure mode i , based on times when the failure mode is not *active*

$\mathbf{t}_{nom,k}$ – [3-by-1 vector] nominal torque on the vehicle about the nominal center of mass (CM) due to the thrusters firing at time step k

3.2 Cataloging failure modes

$\mathbf{a}_{disturbing,i}$ is pre-calculated for every possible failure mode. Multiple-jet failure modes require further cataloging of each combination of thrusters that may be *active*. This

cataloging is done pre-flight, and the values are updated periodically based on the state of the blowdown (the nominal strength of all thrusters drops as the tanks empty).

3.3 Estimating angular acceleration

$\hat{\mathbf{a}}$ is calculated at each FDI update based on the previous 5 gyro samples (covering one full control interval). Assuming small angular rates (so axes are dynamically de-coupled) and that acceleration is constant during each control time period (corresponding to thruster firing times), the acceleration is estimated by fitting a line to the data and taking the slope as shown by the solid line in Figure 3. This least-squares fit is implemented as a computationally efficient linear FIR filter. When performing the fit, the line segment is constrained to begin where the previous segment ended, leading to a contiguous line and improving the estimate. A different $\hat{\mathbf{a}}$ estimation algorithm or sensor may be used with no changes to the rest of the FDI algorithm.

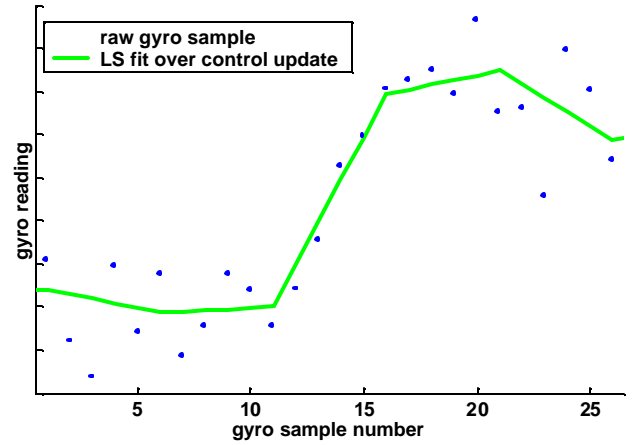


Figure 3: Estimation of angular acceleration

3.4 Calculating nominal-system acceleration

$\mathbf{a}_{nom-system}$ is calculated assuming no failures are present and all physical parameters are at their nominal value (identified values can be used as an alternative, as mentioned in Section 3.11). The force from each thruster, $F_{nom,k}$, resulting torque from each thruster, $\mathbf{t}_{nom,k}$, and finally the vehicle equations of motion from Section 2 (with $\mathbf{w}_{meas,k}$ coming directly from the gyros) are used as follows.

$$\mathbf{F}_{nom,k} = \Gamma_{blowdown} \mathbf{F}_{nom} T_{command,k}$$

$$\mathbf{t}_{nom,k} = (L \times D) \mathbf{F}_{nom,k}$$

where the $L \times D$ cross-product is taken on each column.

$$\mathbf{a}_{nom-system,k} = I_{nom}^{-1} (\mathbf{t}_{nom,k} - \mathbf{w}_{meas,k} \times I_{nom} \mathbf{w}_{meas,k})$$

The disturbing acceleration can then be calculated.

$$\hat{\mathbf{a}}_{disturbing} = \hat{\mathbf{a}} - \mathbf{a}_{nom-system}$$

3.5 Windowing

If the signal-to-noise ratio were high enough, maximum likelihood FDI analysis of the $\hat{\mathbf{a}}_{disturbing}$ readings could be carried out on the values at each time step, as was done in [2]. However, in this application, sensor noise and mass property variations require that values from multiple time steps be combined. Since it is known that failures will occur abruptly, a windowing method is preferred over an IIR (e.g., exponential) filter that would carry through information for longer. In this application a window size of 10 (equal to one second) was found to provide a good balance between speed of response and accuracy. Also, a minimum of 5 samples is required before maximum likelihood FDI analysis is allowed to proceed for a given failure mode.

3.6 Collecting measurements for individual failure modes

As mentioned earlier, one of the challenges of FDI for systems with on-off actuators is that failures are only observable when *active* (as defined in Section 3.1). For example, “off” failures are observable only when the jets are commanded to fire. For each failure mode, only the relevant $\hat{\mathbf{a}}_{disturbing}$ measurements are stored. So for failure mode #1, any time RCS jet 1 is commanded to fire, the resulting $\hat{\mathbf{a}}_{disturbing}$ is logged. These two steps of windowing and collecting data can be considered a type of filtering; however implementation as described here avoids introducing any phase lag between the cause (thruster firings) and effect (vehicle motions), as would be introduced by a linear IIR or Kalman Filter, that would bias the FDI.

3.7 Maximum likelihood

Although the acceleration estimator is nonlinear and sub-optimal, it is reasonable to assume that the estimated disturbing acceleration readings, $\hat{\mathbf{a}}_{disturbing}$, are normally distributed about the true disturbing acceleration values, $\mathbf{a}_{disturbing^*}$. So the probability density for the true disturbing acceleration values, $\mathbf{a}_{disturbing^*}$, conditioned on the measurement history M, is [2] [3]

$$p(\mathbf{a}_{disturbing^*} | M) = (2\pi)^{-3/2} |P_a|^{-1/2} \dots e^{-\frac{1}{2}(\mathbf{a}_{disturbing^*} - \hat{\mathbf{a}}_{disturbing})^T P_a^{-1} (\mathbf{a}_{disturbing^*} - \hat{\mathbf{a}}_{disturbing})}$$

Given disturbing acceleration measurements, $\hat{\mathbf{a}}_{disturbing}$, and knowing the disturbing acceleration values corresponding to each possible failure mode, $\mathbf{a}_{disturbing,i}$, the most likely failure mode is found by finding the $\mathbf{a}_{disturbing,i}$ that maximizes this probability density function. The subscript i indicates the failure mode number

corresponding to the disturbing acceleration. This function is maximized when the likelihood argument, $\lambda_{active,i}$, in the following expression is minimized:

$$I_{active,i} = (\mathbf{a}_{disturbing} - \hat{\mathbf{a}}_{disturbing})^T P_a^{-1} (\mathbf{a}_{disturbing} - \hat{\mathbf{a}}_{disturbing})$$

This expression is calculated and used both to detect and to identify failures. The likelihood argument, $\lambda_{inactive,i}$, is also calculated, using the same equation as above, but using data from periods where the failure mode was not *active*.

3.8 Fault detection

At each FDI update, for each possible failure mode, $\lambda_{active,i}$ is evaluated using the windowed readings. The likelihood argument corresponding to no failure, $\lambda_{active,i0}$, is evaluated using the same windowed relevant readings, but with zero substituted for $\mathbf{a}_{disturbing,i}$. A fault is detected when the ratio of likelihood arguments, $\lambda_{active,i}/\lambda_{active,i0}$ falls below a threshold; this is a generalized likelihood ratio test [8]. Further tests are then performed before identifying a particular failure mode, as described below. Evaluation of individual $\lambda_{active,i}$'s for each failure mode is critically important – evaluating $\lambda_{active,0}$ based on all (windowed) data may not indicate a failure if a failed-off thruster has not fired recently.

3.9 Fault identification

After a fault has been detected, at each FDI update, the likelihood arguments, $\lambda_{active,i}$ and $\lambda_{inactive,i}$, are calculated using all relevant data since the time of detection and compared to certain thresholds and to each other. If a failure mode is true, both $\lambda_{active,i}$ and $\lambda_{inactive,i}$ should be low, indicating the failure mode i fits the data well when it is both *active* and *inactive*. This is used first to remove failure modes from consideration – if $\lambda_{active,i}$ or $\lambda_{inactive,i}$ ever rise above a threshold, failure mode i is decided to be false and removed from further consideration. Then, for a fault to be identified, $\lambda_{active,i}$ must be below a “low” threshold while no other faults are below a high threshold.

Some faults are virtually indistinguishable from one another (in terms of the resulting $\hat{\mathbf{a}}_{disturbing}$), such as this set of four (referring to Figure 2): axial 1 off, axial 2 off, axial 5 on, axial 6 on. The on vs. off failure modes could be distinguished if translational accelerations were used for FDI. The alternative approach taken here to identify the failed thruster is to alter the axial firing pattern (e.g., changing from 1-2-3-5-6-7 on to 2-3-4-6-7-8 on) while maintaining symmetry. Since the firing pattern is adjusted to identify the failed thruster, once the failure has been identified, the pattern is left in a state that makes the failure *inactive*, providing reconfiguration as well as FDI in this case. FDI-driven excitation of failure modes such as this example is generally valuable in expediting the identification.

3.10 FDI based on RLS analysis

In an initial attempt at solving the FDI problem for the X-38, the authors used recursive least squares (RLS) analysis. As had been done in [10], thruster parameters were identified using an exponentially weighted RLS algorithm. This approach did not provide sufficiently reliable FDI for the X-38 application for three main reasons:

1. Relatively high noise levels were present (primarily due to gyro noise and pulse-to-pulse thruster variation).
2. Exponential weighting meant that thrusters fired relatively sparsely (e.g., RCS thrusters as compared to axial thrusters) were not identified well
3. Since multiple axial thrusters are fired continuously, observability of those parameters was very low.

A second, “targeted” RLS-based approach used multiple RLS algorithms, each one identifying the strength of a single thruster with the assumption that all other thrusters were operating nominally. This effectively addressed problems 2 and 3 above, but problem 1 remained. Also, the assumption that all other thrusters are nominal causes partial false positives when the failed thruster fires at the same time a good thruster fires. Methods were developed to address these remaining problems, but results were not sufficiently reliable, motivating development of the maximum-likelihood-based solution.

3.11 Efficiency, Extensions

Many of the terms needed in this analysis, such as $\mathbf{a}_{disturbing,i}$, can be pre-computed or updated periodically.

The algorithm is then relatively efficient. It scales better than linearly as more failure modes are added, since some information is shared between analyses of different failure modes (e.g., estimating $\hat{\mathbf{a}}_{disturbing}$).

This method extends naturally to include translational as well as angular accelerations. This has been implemented in simulation and provides better discrimination between faults since the comparison space is of higher dimension. It was not included in the results presented here since the gyros provided sufficient performance for the X-38 application, and to demonstrate that the method will work for systems with gyros only.

This algorithm has been applied successfully to two other vehicles, the Mini-AERCam mentioned in Section 1, and a 3-dof air-bearing vehicle at the NASA Ames SSRL, demonstrating its generic applicability.

Since it is calculated using nominal mass properties (center of mass location and I_{nom}), the $\hat{\mathbf{a}}_{disturbing}$ estimate is sensitive to off-nominal mass properties. An RLS-based mass property ID method has been developed and implemented in simulation to address this issue, although it was not used to generate these results.

4. FDI applied to the X-38

The FDI algorithm was applied to the X-38, with 40 different failure modes simulated, including each of the 8 RCS and 8 axial thrusters being failed-off or failed-on (32 single-jet failures) and 4 pairs of RCS jets being failed off or on (8 multiple-jet failures). Every mode has been tested multiple times and detection and identification is always accurate and within 5 seconds. Fault detection usually takes only 0.5 seconds, and most failures are identified within about 1.0 second¹. The switching of axial thrusters to distinguish between similar failure modes in some cases causes the time for identification to approach 5 seconds.

An example case is discussed here and shown in Figure 4, for RCS jet 1 failed off. The top part of Figure 4 shows the thruster firing history during this 33-second run. The first 8 rows show the RCS jets pulsing to regulate attitude. The next 8 rows show that axial jets 2-3-4-6-7-8 were on continuously during this run. The next 4 rows correspond to the multiple-jet failure cases, and show when at least one of the jets was commanded to fire. Below that is a zoomed in view of the detection and identification of RCS jet 1 failure. Below that is a legend corresponding to the thruster history as well as the animation screen below. The bottom part of the figure shows a rear view of the vehicle with thrusters firing, torque monitors indicating the net torque produced by the axial and RCS thrusters, and the fault identification result along with a visualization of the likelihood argument, $\lambda_{active,i}$ by drawing a rectangle with width $exp(-0.5\lambda_{active,i})$.

In this simulation, the vehicle starts off with initial angle and rate errors that are largely corrected by thruster firings in the first two seconds. RCS jet 1 abruptly fails off at 3 seconds, indicated by the gray rectangle, but it is not detected until after it fires 23 seconds later. The fault is detected at 26.5 seconds (indicated by the vertical red line), after 0.6 seconds of firing, and is identified at 29.7 seconds (indicated by the change in color from green to red), after a total of 1.0 seconds of firing.

The animation screen at the bottom of Figure 4 was from the final update of this simulation run. RCS jets 1 and 6 are both commanded to fire (as also seen in the thruster history screen), but RCS jet 1 is drawn red, indicating that it has failed. The axial-thruster torque monitor shows minimal torque since the axial thrusters are fired symmetrically and the CM is near the center of the jets. The RCS-thruster torque monitor shows a yaw and a roll torque, caused by RCS jet 6.

¹ Since failure modes may not be observable depending upon whether their thrusters are commanded to fire, the detection and identification times listed indicate the total duration for which the failure was *active*.




Figure 4: Example simulation run for the X-38

The likelihood monitor bars on the right side are drawn with width equal to $\exp(-0.5\lambda_{active,i})$ so they approach 1.0 if the failure is true. This value is close to 1.0 for failure mode 1, and since RCS jet 1 has been identified as failed, it is highlighted in red. RCS jet 2 failed on produces a disturbing acceleration signature close to that of RCS jet 1 failed off, which is why $\exp(-0.5\lambda_{active,22})$ reads above zero (about 0.25). The situation is similar for RCS jet 4 failed on.

In extended testing, the FDI system presented here correctly identified failures in 99.98% of test cases.

5. Conclusions

A maximum-likelihood-based thruster FDI algorithm has been developed and applied in simulation to the X-38 spacecraft. The algorithm is capable of reliably detecting and identifying hard, abrupt single- and multiple-jet on- or off-failures within 1-5 seconds. The algorithm as presented uses gyro signals only, making it applicable to a large number of spacecraft; however, extension to additionally use accelerometer signals has since been implemented, providing even better discrimination between similar failures. The algorithm is computationally efficient and scales better than linearly with the number of failure modes to be identified.

Acknowledgements

The authors wish to thank Rodolfo Gonzalez, Dr. Steven Fredrickson, and Tim Straube of NASA Johnson Space Center and Dave Hammen of LinCom Corp. for providing valuable insights and information for the X-38 and Mini-AERCAM FDI applications. They also wish to thank William Readdy and Gary Martin of NASA headquarters for programmatic support.

References

- [1] Chobotov, *Spacecraft Attitude Dynamics and Control*, 1991, p.14
- [2] Deyst, J. J. and Deckert, J. C. Maximum likelihood failure detection techniques applied to the Shuttle RCS jets. *Journal of Spacecraft and Rockets*, vol. 13, no. 2, 65-74, February 1976
- [3] Gelb, A, *et al. Applied Optimal Estimation*, MIT Press, Cambridge, Mass., 1974.
- [4] Isermann, R. Process Fault Detection Based on Modeling and Estimation Methods – A Survey, *Automatica*, Vol. 20, No. 4, pp. 387-404, 1984.
- [5] Lee, A.Y. and Brown, M.J. A model-based thruster leakage monitor for the Cassini spacecraft. In *Proceedings of the American Control Conference*, 1998, vol. 2, pp 902–904, 1998.
- [6] MATLAB is a registered trademark of The MathWorks, Inc., 24 Prime Park Way, Natick, MA 01760, 508-647-7000.
- [7] NASA Figure.
- [8] Van Trees, H.L. *Detection, Estimation, and Modulation Theory, Part I*, Wiley, New York, 1968.
- [9] Willms, B. Space integrated GPS/INS (SIGI) navigation system for Space Shuttle. In *Digital Avionics Systems Conference Proceedings*, 1999.
- [10] Wilson, E. and Rock, S.M. Reconfigurable control of a free-flying space robot using neural networks. In *Proceedings of the American Control Conference*, Seattle WA, June 1995.
- [11] Wilson, E. *Experiments in Neural Network Control of a Free-Flying Space Robot*. PhD thesis, Stanford University, Stanford, CA 94305, March 1995.
- [12] <http://spaceflight.nasa.gov/station/assembly/elements/x38/>