Proceedings of the 2004 IEEE
International Conference on Robotics & Automation
New Orleans, LA • April 2004

State, Shape, and Parameter Estimation of Space Objects from Range Images

Matthew D. Lichter and Steven Dubowsky
Department of Mechanical Engineering, Massachusetts Institute of Technology
Cambridge, MA 02139 USA
{ lichter | dubowsky }@mit.edu

***Abstract* – An architecture for the estimation of dynamic state, geometric shape, and model parameters of objects in orbit using on-orbit cooperative 3-D vision sensors is presented. This has application in many current and projected space missions, such as automated satellite capture and servicing, debris capture and mitigation, and large space structure assembly and maintenance. The method presented here consists of three parts: (1) kinematic data fusion, which condenses sensory data into coarse kinematic surrogate measurements; (2) Kalman filtering, which filters these surrogate measurements and extracts the full dynamic state and model parameters of the target; and (3) shape estimation, which uses filtered pose information and the raw sensory data to build a body-fixed probabilistic map of the target's shape. This method does not rely on feature detection, optical flow, or model matching, but rather exploits the well-modeled dynamics of objects in space using the Kalman filter. The architecture is computationally fast since only coarse measurements need to be provided to the Kalman filter. This paper will illustrate the three steps of the architecture in the context of rigid body (satellite and debris) estimation and flexible structure estimation.**

Keywords – robot vision; state estimation; shape estimation; parameter estimation; Kalman filter; range images; satellite capture; space debris; space structures; cooperative sensing

I. INTRODUCTION

Many current and future on-orbit space operations will involve automated physical interaction with dynamic free-floating and free-flying targets. Examples include the capture and servicing of satellites, the capture and disposal of space debris, and the robotic assembly and maintenance of large space structures [1-3]. For these missions, it is critical to have accurate knowledge of the target's motions, shape, and dynamic model parameters so that interactions with the target can be planned accordingly. This information is usually unavailable from earth-based sensors and a priori information regarding the shape of the target is often uncertain, especially in the case of damaged satellites, space debris, or thermally deformed or vibrating structures. Therefore, current mission concepts are exploring the use of on-orbit 3-D vision sensors to estimate this information [4].

This paper presents a method for simultaneously estimating dynamic state, geometric shape, and mass model parameters for an arbitrary dynamic target, using only sequences of range images generated by a team of cooperating sensors (see Fig. 1). It is assumed that the sensors' relative positions and orientations are accurately known, and that the target is within the field of view of each of the sensors. A priori knowledge of the properties or geometry of the target is not required, although when available it can be used to speed estimation, reduce sensory requirements, or detect target anomalies (e.g. structural damage). This non-specific need for a priori information makes it well-suited to the applications described above.

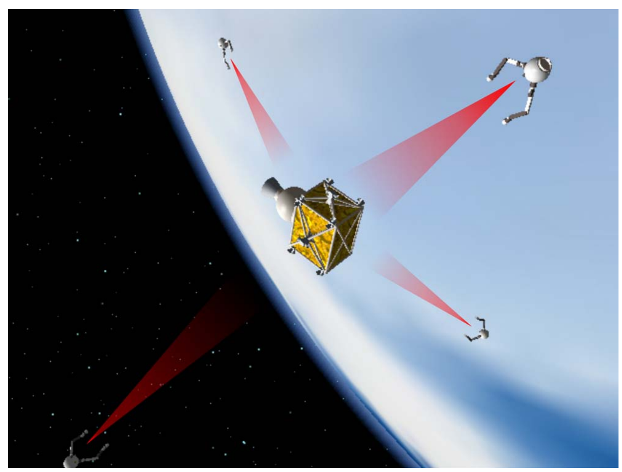

Figure 1. Cooperative vision sensors estimating satellite dynamics.

II. BACKGROUND LITERATURE

The fundamental problem at hand is the simultaneous estimation of both the structure of an object and its motion relative to the observers. While many researchers have studied motion estimation from known shape (and vice versa), simultaneous estimation of both is far more difficult. Notable solution approaches are described below.

A. Feature-based Methods

Methods have been proposed which rely on the continuous tracking of high-level features to determine relative motions of an unknown object or environment. By maintaining an inventory of located features, the methods also estimate high-level geometric structure. Typically a Kalman filter is used to estimate both the feature locations and the motion parameters in a joint framework. Feature-based methods have examined

This work is sponsored by the Japan Aerospace Exploration Agency (JAXA).

the estimation of an unknown object moving with respect to a fixed observer. Natural features of the object (e.g. corners, edges, markers) are extracted and tracked in time to understand the high-level motions and the general structure of the rigid target. References [5] and [6] appear to be the first to employ the Kalman filter for efficient recursive estimation and to employ a mechanics-based physical model of object motions.

Perhaps the most widespread use of feature-based methods in recent years has been in the area of Simultaneous Localization and Mapping (SLAM) [7,8]. Here the task is to construct a feature map of a (typically) static environment while constantly localizing the moving sensors with respect to this map. This is a slight reformulation of the previous problem; the camera is moving instead of the target object.

For space applications, however, these methods are not robust because they are dependent upon feature detection. In practice, many phenomena including occlusions, harsh lighting, and reflective materials can make the reliable detection and correspondence of features virtually impossible. Additionally, these methods by themselves do not provide a detailed estimate of the shape of the target; they provide only a sparse set of feature points pertaining to the object. For this reason, they do not fully address the estimation requirements of the space applications discussed here.

B. Pixel-level Methods

At the opposite end of the spectrum of methodologies are those which rely on pixel-level information rather than high-level features. Several methods use shape from shading, shape from texture, optical flow, or some combination thereof to compute an estimate of object shape at each time step from monocular cameras. Other methods obtain shape estimates directly from stereo cameras or laser technologies. These shape estimates are effectively mosaicked at a pixel level over time, with relative camera motions estimated between time steps [9-11]. The recursive estimation is typically performed via a Kalman filter.

Because of their pixel-level computations, however, these methods are not well-suited to space applications. First, shape computations are highly sensitive to pixel-level detail, which is easily corrupted by the harsh lighting, reflective materials, and highly convoluted surfaces found in space (e.g. wrinkled metallic films used for thermal protection on satellites). Second, pixel information is estimated directly in the Kalman state, leading to a very high-dimensional filter implementation. This is computationally intensive and not feasible for the applications here.

This work attempts to develop an architecture that is not reliant on feature detection schemes or detailed pixel level computations, yet takes advantage of the well-modeled dynamics found in space and operates in a computationally efficient manner.

III. Estimation Architecture Overview

The estimation method presented here consists of three distinct parts, the details of which can be designed independently. The first part consists of kinematic data fusion, a data reduction process which condenses detailed range image data into coarse kinematic information (surrogate measurements) at each sample time (see Fig. 2). The second part is a Kalman filter, which observes the sequence of surrogate measurements over time and extracts the full dynamic state and model parameters of the target using its dynamic model. The final part is a shape estimator, which reinterprets the raw sensory data using filtered kinematic information to build a body-centered probabilistic map of the target's shape.

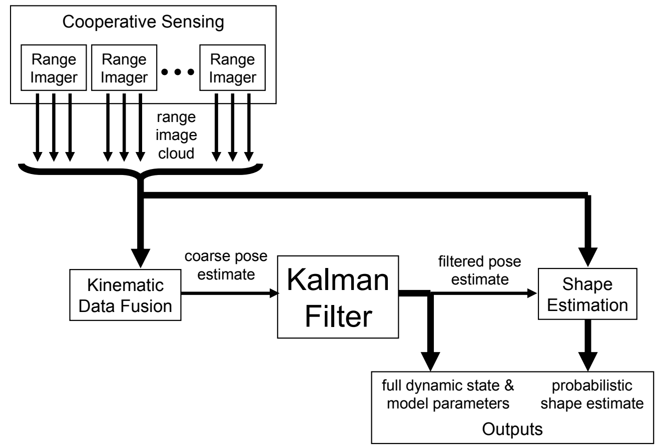

Figure 2. Estimation architecture block diagram.

This approach is powerful in that it takes advantage of the well-modeled dynamics found in space. It does not rely on feature detection schemes, optical flow, or model matching. For this reason, it is more robust to the harsh sensing conditions of space and the fundamental lack of a priori information. Because only rough surrogate measurements are needed for the Kalman filter, the data fusion step can trade accuracy for computational simplicity and speed. Thus, the overall architecture can be very fast. Finally, using sensor uncertainty models while understanding the underlying target motions allows the sensors gain multiple vantage points and fuse redundant noisy measurements in a statistically optimal way.

IV. Satellite and Debris Estimation

For space missions involving the capture of damaged satellites and debris, it is critical to know the translational and rotational positions, velocities, inertial parameters, and geometric shape of the target. This section will illustrate the estimation architecture in the context of this problem.

A. Kinematic Data Fusion

As shown in Fig. 2, the first part of the estimation architecture must reduce range image data into coarse kinematic information of the target at each sample time. This step is effectively a data reduction process, condensing detailed visual data into target pose information that can be digested easily by the Kalman filter (Section IV.B). By the strength of the Kalman filter and the accurate dynamic model, one can afford relatively coarse estimates of pose at this point; therefore this step should focus on computational simplicity and robustness in its processing of the range image data.

Sensory data is assumed to come from a synchronized team of range imagers such as stereo cameras or laser-based technologies. The relative kinematics between the sensors is assumed to be known to high accuracy to allow the mosaicking of data amongst sensors. If several sensors are distributed fairly uniformly about the target and the target is within their field of view, then the surface geometry of the target should be captured reasonably well to first order. The sensors' view of the target will be a noisy 3-D cloud of points as shown in Fig. 3b.

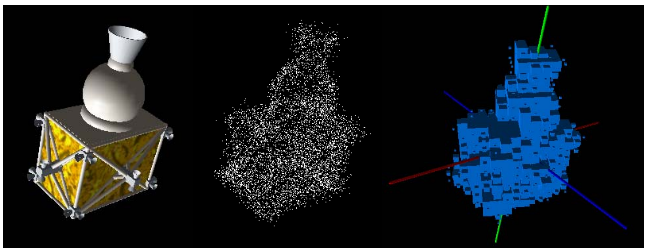

(a) (b) (c)

Figure 3. (a) Simulated object; (b) range image cloud seen by sensors; (c) voxelized representation of cloud, with principal geometric axes shown.

Finding the geometric centroid of this cloud could provide a rough estimate of target position in space. However, this is not robust since sensors which are close to the target will provide a higher density and larger number of sample points than sensors located farther away. Thus the centroid computations would be biased towards closer sensors. A better method is to discretize the space into voxels (volume elements), with each voxel having an occupancy level proportional to the number of sample points found within it (see Fig. 3c). The occupancy values can then be saturated at some pre-defined threshold to minimize the bias effects of closer sensors.

If the target is rigid, this voxel representation is approximately constant in shape, and thus tracking the centroid and principal geometric axes of the voxel image provides a simple way to coarsely track the target's pose. These quantities can be computed in an analogous manner to computing the center of mass and principal inertial axes of a solid body [12]. The centroid position (denoted $\vec{r}_m$) and the rotation matrix describing the attitude of the principal axes (denoted $[R_m]$) thus represent the output of the kinematic data fusion step; they are the surrogate measurements of target position and attitude which feed into the Kalman filter. It is important to note that these tracked axes do not correspond to the actual center of mass or principal inertial axes of the target; they are based merely on superficial geometry of the body.

The computationally simple method presented here has important degeneracies, namely when the target has a high degree of axial symmetry. However, this degeneracy can be handled fairly easily in practice (see [12]). It should be noted that there are many possible ways to provide a surrogate measurement to the Kalman filter. The method presented here is computationally fast, easy to implement, and appears to be quite adequate in many situations. It is used here as an illustrative example to demonstrate the concept of the estimation architecture as a whole.

B. *Kalman Filtering*

The Kalman filter forms the core of the estimation architecture, using the surrogate measurements along with an accurate dynamic model to extract the full dynamic state and inertial parameters of the target. The dynamic state consists of rotational and translational positions and velocities. External forces and torques on the target are assumed to be negligible. Gravity gradient torques and orbital mechanics effects can be incorporated into the model; however their contribution is negligible over short time intervals and in practice it is usually sufficient to model them as process noise in the Kalman filter.

The parameters to be estimated include the principal inertias of the target (relative magnitudes only) and the kinematics of the principal inertial axes and center of mass with respect to the principal geometric axes (surrogate measurement). As mentioned above, the reference frame observed by the data fusion portion does not correspond to the principal inertial axes of the target (see Fig. 4). However, since both sets of axes are fixed to the body, the relative kinematics between the frames is constant and can be parameterized and estimated by the Kalman filter.

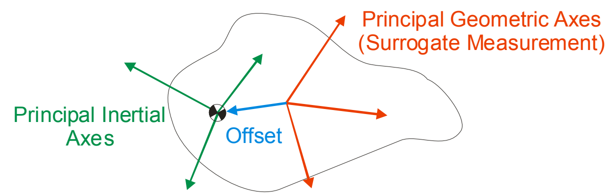

Figure 4. Body-fixed reference frames on the target.

1) *Rotational estimation*

Because the rotational and translational dynamics of the target are decoupled, the estimation can be performed using two separate Kalman filters. For rotational estimation, the target object is assumed to be rigid and external torques on the body are assumed to be negligible, which is a reasonable model for most situations. Rotations are represented using unit quaternions (Euler parameters).

To be estimated are the rotational velocities in body coordinates ($\vec{\omega}_b$), the attitude quaternion of the principal inertial axes ($\vec{q}_b$), the principal inertias of the target ($\vec{I}$), and the rotational offset between the surrogate measurement and the principal inertial axes ($\vec{q}_d$). Only relative magnitudes of the inertias are observable, so $\vec{I}$ is normalized. These quantities follow the deterministic model

$$\frac{d}{dt}\begin{Bmatrix}\omega_{b1}\\ \omega_{b2}\\ \omega_{b3}\\ q_{b0}\\ q_{b1}\\ q_{b2}\\ q_{b3}\\ \vec{I}\\ \vec{q}_d\end{Bmatrix}=\begin{Bmatrix}\frac{I_2-I_3}{I_1}\omega_{b2}\omega_{b3}\\ \frac{I_3-I_1}{I_2}\omega_{b3}\omega_{b1}\\ \frac{I_1-I_2}{I_3}\omega_{b1}\omega_{b2}\\ \frac{1}{2}\left(-\omega_{b1}q_{b1}-\omega_{b2}q_{b2}-\omega_{b3}q_{b3}\right)\\ \frac{1}{2}\left(\omega_{b1}q_{b0}-\omega_{b2}q_{b3}+\omega_{b3}q_{b2}\right)\\ \frac{1}{2}\left(\omega_{b1}q_{b3}+\omega_{b2}q_{b0}-\omega_{b3}q_{b1}\right)\\ \frac{1}{2}\left(-\omega_{b1}q_{b2}+\omega_{b2}q_{b1}+\omega_{b3}q_{b0}\right)\\ \vec{0}\\ \vec{0}\end{Bmatrix} \quad (1)$$

This represents the dynamics of a rigid body under torque-free motion [13] and specifies the parameters to be estimated as constants. The relationship between the surrogate measurement and the principal axes is given by

$$\vec{q}_m = \vec{q}_b \otimes \vec{q}_d \otimes \vec{q}_n \quad (2)$$

where $\vec{q}_m$ is the quaternion parameterization of the surrogate measurement $[R_m]$, $\vec{q}_n$ is the noise quaternion on the surrogate measurement, and the operator $\otimes$ is used to signify quaternion multiplication [13]. Equation (2) merely states that a constant rotation exists between the surrogate measurement and the principal inertial axes of the body.

Equations (1) and (2) represent the process and measurement models for the rotational Kalman filter. These equations present several challenges: they are highly nonlinear; the quaternions and inertia vectors contain normality constraints; and the sum of any two inertia components must be greater than the third (by natural physical bounds). Further, the state estimation errors are kinematically correlated; for example, errors in the estimation of the principal inertia axes' attitude are not independent of errors in the estimation of the offset quaternion.

The first two challenges can be addressed elegantly using the method of Crassidis and Markley [14]. The remaining challenges can be handled through a re-parameterization of the inertia vector to an unbounded space, and coordinate changes to kinematically de-correlate estimation errors in the state vector. Further discussion is beyond the scope of this paper.

2) Translational estimation

Translational estimation is much simpler than rotational estimation. To be estimated is the translational velocity ($\vec{v}_b$) and position ($\vec{r}_b$) of the center of mass, and the translational offset between the surrogate measurement and the center of mass ($\vec{r}_d$) (in a body-fixed coordinate frame). In the absence of external forces, the process and measurement models are given by

$$\frac{d}{dt}\vec{v}_b = \vec{0} \qquad \frac{d}{dt}\vec{r}_b = \vec{v}_b \qquad \frac{d}{dt}\vec{r}_d = \vec{0} \quad (3)$$

$$\vec{r}_m = \vec{r}_b + \left[R_m\right]\cdot\vec{r}_d + \vec{r}_n \quad (4)$$

where $\vec{r}_n$ is positional noise on the surrogate measurement and $[R_m]$ is the rotation matrix from the surrogate measurement. It is usually more appropriate to use a filtered version of $[R_m]$ from the rotational filter rather than using the much noisier surrogate measurement directly. While this does introduce slight coupling between the filters, the performance tends to be superior in practice. Since (3) and (4) are linear, only a basic (linear) Kalman filter [15] is needed.

C. Shape Estimation

With accurate knowledge about the trajectory of the target, shape estimation reduces to a classic stochastic map-building problem. Since the environment (target) motions are known with respect to the sensors, pixel-level data can be fused into a probabilistic map of the target's shape using appropriate sensor noise models. Numerous methods can be used to do this [16,17]. One simplistic approach will be illustrated here.

At each sample time, a probability density function (PDF) in 3-D space can be generated to statistically describe the relative likelihood that the target exists at a given point in space. This overall PDF can be obtained by convolving each sensor data point with the PDF describing the noise characteristics of the sensor. That is

$$p_{shape}\left(\vec{r}\right) = \sum_i \delta\left(\vec{r}-\vec{r}_i\right) * p_{sensor}\left(\vec{r}\right) \quad (5)$$

where $p_{shape}(\vec{r})$ is the PDF describing the relative likelihood that the target surface exists at point $\vec{r}$, $p_{sensor}(\vec{r})$ is the PDF describing the noise distribution at each sensor data point, $\vec{r}_i$ is the location of the i^{th} data point, $\delta(\ )$ is the Dirac delta function, and * is the convolution operator. The goal of shape estimation is to accurately depict and sharpen $p_{shape}(\vec{r})$.

It is important that the coordinate system used in (5) be a body-fixed reference frame, so that data can be fused across time as the target moves. It is best to use a filtered version of the surrogate measurement frame, rather than the principal inertial axes (see Fig. 4), because it can be tracked even under degenerate conditions [12].

Equation (5) can be computed at each sample time to obtain a PDF based on the current data. This should be combined somehow with the PDFs from all previous sample times to yield a more accurate overall PDF. A computationally efficient method would be recursive, making incremental improvements to its estimate using new information. An ad hoc solution would be to compute the new estimate as a weighted average of the previous estimate and the new information. This amounts to applying a forgetting factor to old data:

$$p_{shape}{}^{1\to t} = \alpha \cdot p_{shape}{}^{1\to(t-1)} + \left(1-\alpha\right)\cdot p_{shape}{}^{t} \quad (6)$$

where $p_{shape}{}^{1\to t}$ is the cumulative shape estimate incorporating information from time samples 1 through t, $p_{shape}{}^{1\to(t-1)}$ is the old cumulative estimate, $p_{shape}{}^{t}$ is the PDF based strictly on the current sensory data, and α is a forgetting factor between 0 and 1. The larger the forgetting factor, the slower old information is forgotten. The cumulative estimate can be initialized to a priori knowledge of the target shape, if it exists.

In practice, it is generally much easier to deal with discrete sums than integrals in the convolution process. Therefore, the body-fixed space is discretized into voxels, this time with a much finer resolution than that used in Section IV.A. The finer resolution is enabled because data is being used from all the sample times rather than a single sample; therefore much more information is available. Fig. 5 shows the resultant voxel estimate of a simulated target based on 100 sample times, with voxels having a high likelihood of occupancy shown.

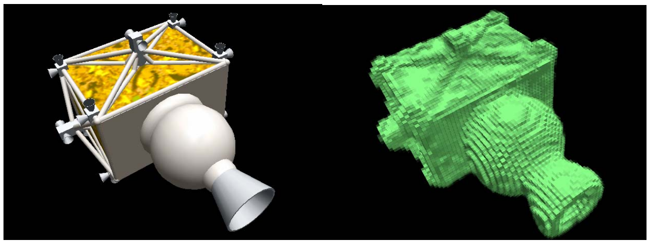

(a) (b)

Figure 5. (a) Simulated target; (b) voxel-based shape estimate.

This is a rather rudimentary approach but demonstrates the concept of the architecture. Clearly, more sophisticated approaches could be applied to this problem as necessary. This is an area of ongoing research.

D. Estimator Performance

Simulation studies suggest that this estimation approach is fast and robust. A virtual environment containing representative targets was constructed and noisy range data was synthesized. The methods described above were used to estimate state, shape, and parameters of these virtual targets. For sufficiently rich trajectories (i.e. targets exhibiting a multi-axis tumble), the state and parameter space was fully observable and the Kalman filter converged rapidly (on the order of one target rotation period, see Fig. 6). While no explicit studies were made on computational requirements, the majority of the burden resulted from data synthesis and graphics processing, and the estimator still operated at a rate of several Hz on an Intel Pentium III processor. Further details on these simulation studies can be found in [12].

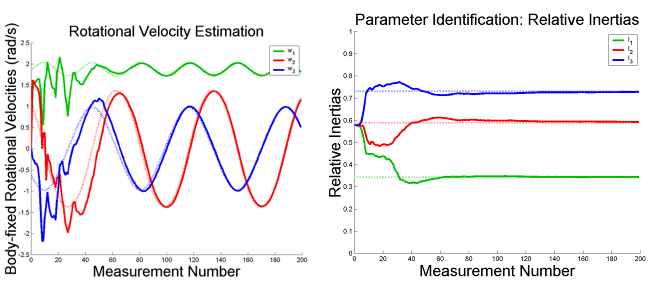

Figure 6. Kalman filter convergence plots from a simulated estimation.

V. LARGE FLEXIBLE SPACE STRUCTURE ESTIMATION

Large space structures tend to exhibit lightly damped, multimodal oscillations in orbit, induced by variable thermal loading and other disturbances (see Fig. 7). In performing tasks such as assembly, maintenance, and inspection of such structures, robotic systems could benefit greatly if they were able to understand and predict these dynamics [18].

While this estimation problem might seem quite different from the previous problem, it can in fact be handled using the same broad approach. This section will describe the application of the general estimation methodology to this class of problem.

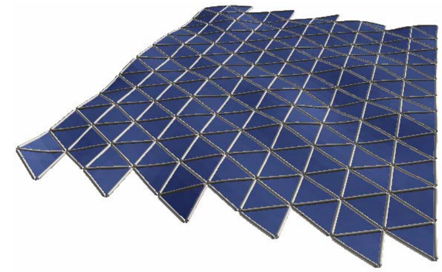

Figure 7. Large space structure undergoing vibration.

It is assumed here that the dynamics are linear and that the mode shapes and modal frequencies are approximately known. This is not a necessary condition and is used only to simplify the illustration here.

A. Kinematic Data Fusion

Again, the purpose of the data fusion step is to reduce the detailed range data into a high-level usable form for the Kalman filter. For linear dynamic systems, this amounts to decomposing the data into each of the modal components. At each sample time, an inner (dot) product is taken with the range data and the known mode shapes:

$$\hat{d}_i(t) \equiv \langle RangeData(\vec{x},t), ModeShape_i(\vec{x}) \rangle_{\vec{x}} \tag{7}$$

where $\hat{d}_i(t)$ is the approximate value for the i^{th} modal component, $RangeData(\vec{x},t)$ is the set of all range data points in space and time, $ModeShape_i(\vec{x})$ is the i^{th} known mode shape of the structure, and the $< a\ ,\ b >_x$ operator indicates the inner product of a and b in the spatial domain. This decomposition need only be done on modes relevant to the mission tasks at hand. For example, modes whose frequencies are outside the bandwidth of the robots could be ignored.

B. Kalman Filtering

The quantities $\hat{d}_i(t)$ for the N important modes are the outputs of the kinematic data fusion step. These quantities oscillate sinusoidally in time with the i^{th} modal frequency. However, noise in the range images and artifacts from imperfect spatial sampling leads to noise in the modal components computed in (7). Since these signals are expected to follow a very weakly decaying sinusoid, they can be filtered easily using a Kalman filter (see Fig. 8).

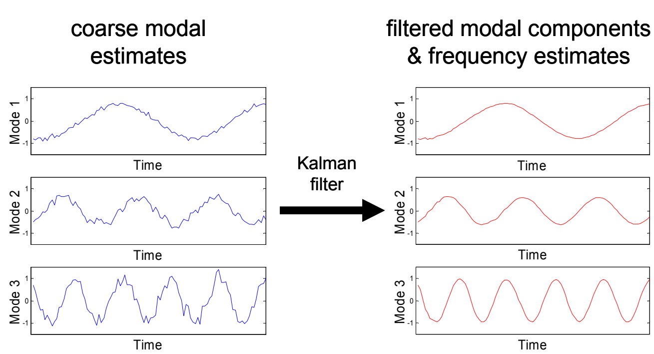

Figure 8. Filtering the surrogate measurements.

The Kalman filter process and measurement models describing each sinusoid are given by the equations

$$\frac{d}{dt}\begin{Bmatrix} d_i \\ v_i \\ \omega_i \end{Bmatrix} = \begin{Bmatrix} v_i \\ -\omega_i^2 d_i \\ 0 \end{Bmatrix} \tag{8}$$

$$\hat{d}_i = d_i + n_i \tag{9}$$

where d_i is the i^{th} modal component, v_i is the time rate of change of that component, ω_i is the fundamental frequency of that mode, and n_i is the noise on the surrogate measurement (7). For illustrative simplicity, the small damping terms have been omitted. If damping were significant, it could be parameterized and estimated by extending (8) accordingly. The fundamental frequencies must be explicitly estimated through the Kalman filter, since they cannot be known perfectly a priori. Since (8) is nonlinear, an extended Kalman filter [15] or an unscented Kalman filter [19] should be used.

C. *Shape Estimation*

For this problem, shape estimation is simply modal reconstruction. That is,

$$ShapeEstimate\left(\vec{x},t\right) = \sum_{i=1}^{N} d_i\left(t\right) \cdot ModeShape_i\left(\vec{x}\right) \tag{10}$$

If desired, range information from the sensors could be incorporated into this shape estimate, in a manner similar to that described in Section IV.C.

VI. Summary

This paper has described a broad estimation approach for the combined estimation of state, shape, and model parameters of space objects from sequences of range images. The approach consists of three parts. A kinematic data fusion step reduces detailed pixel-level information into high-level surrogate measurements that can be processed easily. A Kalman filter exploits the high-fidelity dynamic model of space objects to filter and extract information from these coarse surrogate measurements. A shape estimator then reinterprets raw sensory data from a filtered vantage point and builds a probabilistic map of the target using sensor uncertainty models.

This paper has illustrated the concept for both the estimation of rigid tumbling targets (i.e. satellites and space debris) and large flexible space structures. Although the details for each step of the estimator were different for each application, the overall broad methodology was identical.

Acknowledgments

The authors would like to thank Drs. Noam Brook, Karl Iagnemma, and Vivek Sujan for their helpful dialogue during the development of this work.

References

[1] Bornschlegl, E., G. Hirzinger, M. Maurette, R. Mugnuolo, and G. Visentin, "Space robotics in Europe, a compendium," Proceedings of the 7th International Symposium on Artificial Intelligence, Robotics, and Automation in Space: i-SAIRAS 2003, NARA, Japan, May 19-23, 2003.

[2] Piedboeuf, J-C. and É. Dupuis, "Recent Canadian activities in space automation and robotics – an overview," Proceedings of the 7th International Symposium on Artificial Intelligence, Robotics, and Automation in Space: i-SAIRAS 2003, NARA, Japan, May 19-23, 2003.

[3] Ueno, H., et al, "Space robotic mission concepts for capturing stray objects," Proceedings of the 23rd International Symposium on Space Technology and Science, Matsue, Japan, May 26th-June 2nd, 2002.

[4] Wakabayashi, Y., Y. Ohkami, M. Miyata, T. Adachi, and T. Iijima, "A compact laser range finder for space applications," Proceedings of SPIE, Vol. 3714: Enabling Photonic Technologies for Aerospace Applications, A.R. Pirich, E.W. Taylor, eds., 1999, pp. 131-138.

[5] Broida, T.J. and R. Chellappa, "Estimating the kinematics and structure of a rigid object from a sequence of monocular images," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 13, No. 6, pp. 497-513, 1991.

[6] Young, G-S.J. and R. Chellappa, "3-D motion estimation using a sequence of noisy stereo images: models, estimation, and uniqueness results," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 12, No. 8, pp. 735-759, Aug 1990.

[7] Dissanayake, M.W.M.G., P. Newman, S. Clark, H.F. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localization and map building (SLAM) problem," IEEE Transactions on Robotics and Automation, Vol. 17, Issue 3, pp. 229-241, Jun 2001.

[8] Leonard, J.J. and H.F. Durrant-Whyte, Directed Sonar Sensing for Mobile Robot Navigation, Boston, MA: Kluwer Academic, 1992.

[9] Heel, J., "Direct dynamic motion vision," Proceedings of the IEEE International Conference on Robotics and Automation, Vol. 2, pp. 1142-47, May 13-18, 1990.

[10] Horn, B.K.P. and J.G. Harris, "Rigid body motion from range image sequences," Computer Vision, Graphics, and Image Processing (CVGIP): Image Understanding, Vol. 53, No. 1, Jan, pp. 1-13, 1991.

[11] Soatto, S., R. Frezza, and P. Perona, "Motion estimation via dynamic vision," IEEE Transactions on Automatic Control, Vol. 41, No. 3, March 1996, pp. 393-413.

[12] Lichter, M.D., and S. Dubowsky, "Estimation of state, shape, and inertial parameters of space objects from sequences of range images," Proceedings of the SPIE Conference on Intellligent Robotics and Computer Vision XXI, Vol. 5267, Providence, RI, Oct. 2003.

[13] Sidi, M.J., Spacecraft Dynamics and Control: a Practical Engineering Approach, New York: Cambridge University Press, 1997.

[14] Crassidis, J.L., and F.L. Markley, "Unscented filtering for spacecraft attitude estimation," AIAA Journal of Guidance, Control, and Dynamics, Vol. 26, No. 4, July-August 2003, pp. 536-542.

[15] Brown, R.G., and P.Y.C. Hwang, Introduction to Random Signals and Applied Kalman Filtering, 2nd Ed., New York: John Wiley & Sons, 1992.

[16] Smith, R.C., and P. Cheeseman, "On the representation and estimation of spatial uncertainty," International Journal of Robotics Research, Vol. 5, No. 4, pp. 56-68, 1987.

[17] Sujan, V., and S. Dubowsky, "Visually built task models for robot teams in unstructured environments," Proceedings of 2002 IEEE International Conference on Robotics and Automation (ICRA 2002), Washington, D.C., May 11-May 15, 2002.

[18] Ueno, H., T. Nishimaki, M. Oda, and N. Inaba, "Autonomous cooperative robots for space structure assembly and maintenance," Proceedings of the 7th International Symposium on Artificial Intelligence, Robotics, and Automation in Space: i-SAIRAS 2003, NARA, Japan, May 19-23, 2003.

[19] Julier, S.J., and J.K. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls, Orlando, Florida, 1997.