DECENTRALIZED ESTIMATION ALGORITHMS FOR FORMATION FLYING SPACECRAFT

Philip Ferguson* and Jonathan How†
Massachusetts Institute of Technology

Abstract

Several space missions are being planned with large fleets of spacecraft at GEO and beyond where relative navigation using GPS will either be impossible or insufficient. Onboard local ranging devices will be required to replace or augment the available GPS measurements to perform relative navigation for these missions. The vast amounts of distributed data and the computational load limitations suggest that these navigation algorithms should be highly distributed, but the nonlinear local range measurements couple the states of the vehicles, which greatly complicates this decentralization. This paper investigates several estimator architectures and the associated algorithms for determining the fleet state, and then compares the resulting performance, computation, and communication requirements. This analysis shows that the proposed decentralized reduced-order filters provide near optimal estimation results without excessive communication or computation. Embedding these reduced-order estimators within the hierarchic architecture presented should permit scaling of the relative navigation to very large fleets.

Introduction

The concept of formation flying of satellite clusters has been identified as an enabling technology for many future NASA and the U.S. Air Force missions[1–3] The use of fleets of smaller satellites instead of one monolithic satellite should improve the science return through longer baseline observations, enable faster ground track repeats, and provide a high degree of redundancy and reconfigurability in the event of a single vehicle failure. The GN&C tasks are very complicated for larger fleets because of the size of the associated estimation and control problems and the large volumes of measurement data available. As a result, distributing the guidance and control algorithms becomes a necessity to balance the computational load across the fleet and to manage the inter-spacecraft communication. This is true not only for the optimal planning, coordination and control,[4] but also for the fleet state estimation since the raw measurement data is typically collected in a decentralized manner (*i.e.,* each vehicle takes its own local measurements).

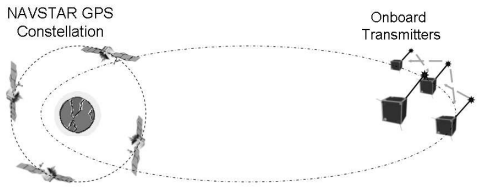

Fig. 1: GPS Estimation with Local Ranging Augmentation. Note: Not to scale.

State estimation is further complicated by nonlinear measurements, requiring the use of extended (often iterated) Kalman filters. A commonly-used, highly accurate sensor for fleet state estimation is the Global Positioning System (GPS). Recent work on GPS estimators for relative navigation in LEO using Carrier-Phase Differential GPS demonstrated 2 cm accuracy in relative position and better than 0.5 mm/s in relative velocity.[5–7] These results were obtained using a fully decentralized filter, and the high accuracy results achieved validate that the relative GPS measurements (single differences relative to the master) taken on one vehicle can be treated as if they are entirely uncorrelated from the single difference measurements taken on other vehicles. Thus the full fleet measurement matrix, H, can be treated as block-diagonal and small coupling effects (such as a differential ionosphere) can be ignored if the fleet separation is less than 10km.[5–7]

While GPS can be used as an effective sensor for many space applications, it requires constant visibility of the GPS constellation. In space, GPS visibility begins to breakdown at high orbital altitudes (e.g. highly elliptic, GEO, or at L2). Thus, a measurement augmentation is desired to permit relative navigation through periods of poor visibility and also to improve the accuracy when the GPS constellation is

* Massachusetts Institute of Technology, philf@mit.edu
† Senior Member AIAA, MIT jhow@mit.edu

visible. Fig. 1 illustrates measurement augmentation using local ranging devices on each vehicle that measure a scalar range and velocity between each pair of vehicles in the fleet.[8–12] However the local range measurements taken onboard the spacecraft strongly correlate the states of the vehicles, which destroys the block-diagonal nature of the full fleet measurement matrix, H, and greatly complicates the process of decentralizing the algorithms.[14] In contrast to the GPS-only estimation scenario which effectively decentralized for reasonably sized fleet separations, this estimation problem does not decorrelate at any level. As a result, methods must be devised to efficiently decentralize the estimation algorithms while retaining as much accuracy as possible. Ultimately, an estimation architecture (both the communication network and computational flow required) is desired that provides accurate relative state estimates in many different estimation regimes. Three basic architectures are studied in this paper:

- **Centralized:** All processing done on one vehicle.
- **Decentralized:** Each vehicle processes its own data.
- **Hierarchic:** Multiple layers of small centralized or decentralized architectures.

To evaluate the viability of these architectures, algorithms must be developed to populate them. This paper presents various estimation algorithms and architectures, developed to mesh appropriately with typical control algorithms and/or communication systems for formation flying missions with large fleets and limited GPS visibility. The following sections present several decentralized estimation techniques and then the FFIT testbed[15] is used to perform a trade study to assess which ones are best suited for certain mission scenarios. As a result of this analysis, a scalable estimation architecture is proposed that conserves fleet communication and computation while achieving near-optimal accuracies for most mission scenarios.

Centralized Architectures

Fig. 2(a) illustrates the basic algorithmic structure of a centralized algorithm on a fleet consisting of three vehicles (for clarity, vehicle number 1 will be designated the *master* and all others as *slaves*). The master gathers all measurement data for processing in a centralized filter. Depending on the fleet control requirements, solutions from the centralized filter may be sent back out to the slave vehicles for control and/or science use. The following sections

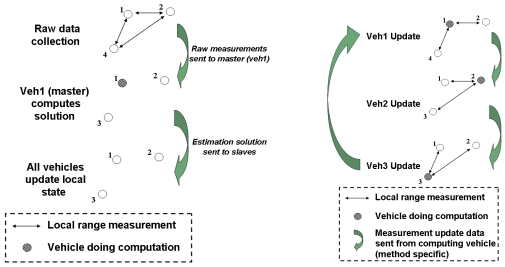

(a) Centralized (b) Cascade

Fig. 2: Basic algorithmic flow for centralized and cascade-type architectures.

present two algebraically equivalent forms of the optimal filter: The Kalman Filter and The Information Filter. Each form has its merits and disadvantages regarding communication and computation, as will be further explored in later sections.

Kalman Filter

The optimal filter that minimizes the mean-squared estimation error of a state vector is the well known Kalman filter given by the following equations:
Kalman Measurement Update

$$S_k = H_k P_k^- H_k^T + R_k \tag{1}$$
$$K_k = P_k^- H_k^T S_k^{-1} \tag{2}$$
$$\hat{X}_k^+ = \hat{X}_k^- + K_k \left(z_k - H_k \hat{X}_k^- \right) \tag{3}$$
$$P_k^+ = \left(\mathbf{1} - K_k H_k \right) P_k^- \tag{4}$$

Kalman Time Update

$$\hat{X}_{k+1}^- = \Phi_k \hat{X}_k^+ \tag{5}$$
$$P_{k+1}^- = \Phi_k P_k^+ \Phi_k^T + Q_k \tag{6}$$

In this *centralized* filter, each slave sends its local measurement vector, z_{k_i}, to the master who executes Eqs. 1–6 at every time-step. For integration with the control architecture, the master may need to broadcast the estimation solutions to the slaves.

Information Filter

The information filter is an algebraically equivalent form of the Kalman filter, cast in a light to make explicit the information content of each measurement and how it impacts the global state estimate. The information filter requires the introduction of the following *information* variables:

$$Y_k^{\pm} \triangleq P_k^{\pm^{-1}} \tag{7}$$

$$\hat{y}_k^{\pm} \triangleq Y_k^{\pm}\hat{X}_k^{\pm} \quad (8)$$

$$i_k \triangleq H_k^T R_k^{-1} z_k \quad (9)$$

$$I_k \triangleq H_k^T R_k^{-1} H_k \quad (10)$$

With these definitions, the entire information filter is summarized below.
Information Measurement Update

$$\hat{y}_k^+ = \hat{y}_k^- + i_k \quad (11)$$

$$Y_k^+ = Y_k^- + I_k \quad (12)$$

Information Time Update

$$M_k = \left(\Phi_k^{-1}\right)^T Y_k^+ \Phi_k^{-1} \quad (13)$$

$$Y_{k+1}^- = \left[I - M_k\left(M_k + Q_k^{-1}\right)^{-1}\right] M_k \quad (14)$$

$$\hat{y}_{k+1}^- = \left[I - M_k\left(M_k + Q_k^{-1}\right)^{-1}\right] \left(\Phi_k^{-1}\right)^T \hat{y}_k^+ \quad (15)$$

Similar to the Kalman filter, the centralized implementation of the information filter requires all information be gathered at the master for processing at every time-step. The slaves send their local information vector, i_{k_i} and matrix contributions, I_{k_i} based on their local measurements, z_{k_i} to the master, who incorporates the new information as follows:

$$\hat{y}_k^+ = \hat{y}_k^- + \sum i_{k_i} \quad (16)$$

$$Y_k^+ = Y_k^- + \sum I_{k_i} \quad (17)$$

The master then executes Eqs. 13–15 to complete the update cycle. In the sections that follow, the functional forms of these two centralized forms will be used to create several decentralized filters.

Decentralized Architectures

For large fleets, it may be desirable to spread the computational effort of the estimation more uniformly across the fleet. Also, since the raw measurements are gathered in a decentralized fashion (*i.e.,* each vehicle measures its own GPS and local ranging information), a decentralized estimator would reduce the need to communicate the measurements. This section describes several viable decentralized estimators, split up into two classes: *Full Order* and *Reduced Order* filters.

Full Order Decentralized Filters

The first class of decentralized filters is *full order*. A fleet running a full order decentralized estimator will have *each* vehicle estimating the *entire* fleet state.

Decentralized Information Filter: Previous decentralized estimation research[16–18] has relied heavily on the use of information filters as a means of assimilating remotely collected data. The typical argument for this preference lies in Eqs. 11 and 12. Notice that the form of these measurement update equations is such that the updated quantity is simply the old quantity plus the new *information* provided by the new measurement. In a decentralized implementation of this filter, each spacecraft would simply sum the new information from every other vehicle.[19] Furthermore, since the information filter is algebraically equivalent to the centralized form of the filter, no information is lost and the best possible estimate (*i.e.,* the centralized estimate) is available on each spacecraft in the fleet.

The decentralized information as described above is an excellent solution for some scenarios. However, in the case of many vehicles estimating the fleet state, the information filter has some substantial short comings:

1. While the measurement update is considerably simpler than the conventional Kalman filter, the propagation step is much more complicated.
2. The information filter deals strictly in information variables that have little significance to the actual problem. To obtain the actual state estimate and its associated covariance, one must solve Eqs. 7 and 8 together to back out the state estimate and covariance matrix.
3. A fully connected network must exist to trade the data from each vehicle to every other vehicle in the fleet. Furthermore, this data can potentially be quite large since an information vector and information matrix (see Eqs. 7 and 8) must be sent at each update*.

Note that it has previously been suggested[13] that it may be possible for each vehicle in the fleet to determine the measurement matrix, H, for all other vehicles thereby eliminating the need to also transmit the information matrix at each time step. But this does not work well for nonlinear filters because a discrepancy arises from the fact that the linearized H matrix must be obtained using the state estimate to generate the information components I and i in Eqs. 9 and 10. If each vehicle has different estimates of the fleet state, small errors (due to second order variations in $h(\hat{X})$) will accumulate through the additive updates of Eqs. 11 and 12.

*Note that only half of the information matrix need be sent since it is symmetric.

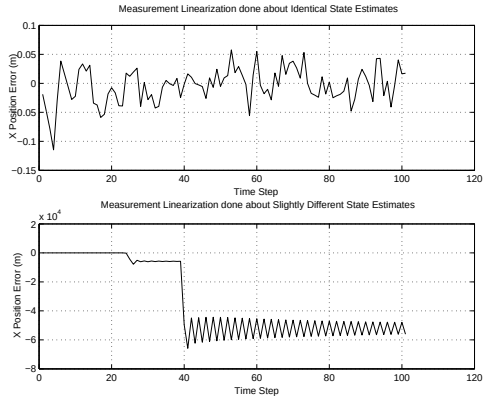

Fig. 3: The effect of measurement equation linearization about slightly different state estimates.

For example, denote the state estimate of the master vehicle as $\hat{X}_1$ and its linearized measurement matrix as $H(\hat{X}_1)$. Assume that the slave vehicle has a slightly different state estimate ($\hat{X}_2 = \hat{X}_1 + \Delta X$), then the linearized measurement matrix for the slave vehicle is

$$H(\hat{X}_2) \approx H(\hat{X}_1) + \left.\frac{\partial H(X)}{\partial X}\right|_{\hat{X}} \Delta X \qquad (18)$$

The second term in Eq. 18 creates a difference in the estimates that can accumulate over time. The result of this second order effect is demonstrated in Fig. 3. In this 2-vehicle simulation, each vehicle computes i_i as per Eq. 9 using an H that was linearized about the slave's current state estimate. The slave then transmits i_k to the master vehicle. To perform the update step, the master must compute the I_k for the slave vehicle based on an H that was linearized about the *master's* current state estimate. In this scenario, the master's state estimate and the slave's state estimate differ by less than a centimeter. Fig. 3 demonstrates the filter going unstable at approximately the 23rd time-step. To prevent this instability, the vehicles must transmit both the information vector *and* the information matrix at every time step (requiring substantially more communication).

Full-Order Iterative Cascade Filter: Another filter considered is the full-order iterative cascade filter. In this algorithm (depicted in Fig. 2(b)) each vehicle employs a standard Kalman filter estimating the full fleet state, but using only the locally available measurements. After the measurement update, each vehicle broadcasts its local state solution to every other vehicle. Upon receipt, the vehicles re-compute their measurement matrices, H_i and re-compute a measurement update. This update procedure is identical to that proposed by Park,[14,20] except that in this case, each vehicle is estimating the entire fleet state, whereas in Park's filter, each vehicle estimated only their local state.

Note that the full-order iterative cascade filter is sub-optimal because there is no single filter in the fleet that can accurately estimate the inter-vehicle correlations (the correlations are approximated by iterations around the fleet). Furthermore, given the amount of computational effort that must be exerted to execute a full-order filter, this method may not provide a good balance between estimation accuracy and computational effort.

Reduced-Order Filters

For many control/science applications, each vehicle is primarily interested in estimating its own state. This prompted the development of *reduced-order decentralized filters* wherein each vehicle estimates only its local state thereby substantially reducing the computational demands on each vehicle at the cost of sub-optimality and increased synchronization requirements. The extent to which the estimator performance is impacted is examined in the Analysis Section. The algorithmic flow for all of the reduced-order decentralized filters is shown in Fig. 2(b).

Cascade Filter: Recent work by Park proposed a reduced-order decentralized filter known as the Iterative Cascade Extended Kalman Filter (ICEKF) (see Refs. [14,20]). The purpose of the ICEKF was to provide a decentralized local ranging augmentation for LEO applications using GPS pseudolites. Furthermore, Park's algorithm was intended for a relatively small number of range measurements (compared to the GPS measurements) with identical accuracy to the other GPS measurements. In this environment, the ICEKF works extremely well and Park has demonstrated near optimal performance.[14,20] For this paper, we are interested not only in local ranging for LEO augmentation, but also for MEO and beyond. For these scenarios, the estimator will have to rely almost solely on the local ranging data that could be much more accurate than the GPS measurements. In this scenario, the ICEKF exhibits poor performance (see Analysis Section for simulation results), typically yielding unstable results.

The primary problem with the ICEKF is that the relative state vectors from the other vehicles are assumed to be perfect, when in reality, these states are in the process of being estimated and thus are erroneous. One *ad hoc* method for accounting for the estimation error associated with other vehicles' states is to absorb it into the measurement error variance matrix R.[24] Before starting the estimator, R is

increased ("bumped up") by a constant amount that corresponds to the other vehicles' estimated error

$$R_{\text{bump}} = R + JP_{\text{yy}}J^T \tag{19}$$

where R is the original measurement error variance matrix, J is the measurement matrix for all other non-local measurements in the fleet and P_{yy} is the initial covariance matrix for all other non-local states in the fleet state vector. Increasing R indicates that the measurements might not be as good as suggested by the accuracy of the ranging device.

Another possible way to incorporate more measurement uncertainty and improve the filter stability/accuracy is to transmit some state covariance information along with the local state vector. The following describes a technique developed to incorporate this state uncertainty information.

Schmidt-Kalman Filter: The traditional purpose of the *Schmidt-Kalman filter* (SKF)[21] is to reduce the computational complexity of a Kalman filter by eliminating states that are of no physical interest, but are required to estimate noises and/or biases. The formulation of the SKF begins with a partitioning of the standard state propagation and measurement equations as well as the covariance matrix:

$$\begin{bmatrix} x \\ y \end{bmatrix}_{k+1} = \begin{bmatrix} \phi_x & 0 \\ 0 & \phi_y \end{bmatrix}_k \begin{bmatrix} x \\ y \end{bmatrix}_k + \begin{bmatrix} w_x \\ w_y \end{bmatrix}_k \tag{20}$$

$$z_k = \begin{bmatrix} H & J \end{bmatrix}_k \begin{bmatrix} x \\ y \end{bmatrix}_k + \nu_k \tag{21}$$

$$P_k = \begin{bmatrix} P_{xx} & P_{yx} \\ P_{xy} & P_{yy} \end{bmatrix}_k \tag{22}$$

where x represents the state vector containing the states of interest and y represents the remaining states. Applying the partitions of Eqs. 20 and 21 to the general Kalman filter equations, solving for each block, and *setting the gain for the y states to zero* yields the following equations:

Schmidt-Kalman Measurement Update[†]:

$$\begin{aligned} \alpha_k &= H_k P^-_{xx_k} H_k^T + H_k P^-_{xy_k} J_k^T + \\ &\quad J_k P^-_{yx_k} H_k^T + J_k P^-_{yy_k} J_k^T + R_k \quad (23) \\ K_k &= (P^-_{xx_k} H_k^T + P^-_{xy_k} J^T)\alpha_k^{-1} \quad (24) \end{aligned}$$

[†] Note that in Eq. 25, the $J_k\hat{y}_0$ term is only required if the y states are expected to have a non-zero mean. If this is the case, a best guess of the $\hat{y}$ states is inserted for $\hat{y}_0$. In many applications, the best guess might be $\hat{y}_0 = 0$.

$$\begin{aligned} \hat{x}^+_k &= \hat{x}^-_k + K_k(z_k - H_k\hat{x}^-_k - J_k\hat{y}_0) \quad (25) \\ P^+_{xx_k} &= (I - K_kH_k)P^-_{xx_k} - K_kJ_kP^-_{yx_k} \quad (26) \\ P^+_{xy_k} &= (I - K_kH_k)P^-_{xy_k} - K_kJ_kP^-_{yy_k} \quad (27) \\ P^+_{yx_k} &= P^{-^T}_{xy_{k+1}} \quad (28) \\ P^+_{yy_k} &= P^-_{yy_k} \quad (29) \end{aligned}$$

Schmidt-Kalman Time Update

$$\begin{aligned} \hat{x}^-_{k+1} &= \phi_{x_k}\hat{x}^+_k \quad (30) \\ P^-_{xx_{k+1}} &= \phi_{x_k}P^+_{xx_k}\phi^T_{x_k} + Q_{x_k} \quad (31) \\ P^-_{xy_{k+1}} &= \phi_{x_k}P^+_{xy_k}\phi^T_{y_k} \quad (32) \\ P^-_{yx_{k+1}} &= P^{-^T}_{xy_{k+1}} \quad (33) \\ P^-_{yy_{k+1}} &= \phi_{y_k}P^+_{yy_k}\phi^T_{y_k} + Q_{y_k} \quad (34) \end{aligned}$$

The SKF equations may appear more complicated than those of the Kalman filter; however, substantial computational savings are embedded in the fact that the filter only solves for the x state and not the y state which, in typical applications of this technique, is of no physical interest. It is this aspect of the SKF that is appealing for the design of reduced-order decentralized filters. In particular, the SKF is used to incorporate a covariance estimate of the relative state estimates, which eliminates the prior assumption that the states of all other vehicles are known perfectly.

This method introduces the relative states of all fleet spacecraft into the state vector, and then uses a Schmidt formulation to incorporate the uncertainty in other vehicles' states, while reducing the estimated state vector to include only the local state. Each vehicle follows the same iteration procedure as the ICEKF, but instead of transmitting just its local state vector to the next vehicle in the fleet, it sends the local state vector along with its local covariance matrix (or some representative portion). Also, instead of executing the standard Kalman filter equations, the SKF equations are used by replacing P_{yy} with the transmitted covariance matrices of each other vehicle placed on a block-diagonal. Thus, with P_{yy} being transmitted to each vehicle at every time-step, Eqs. 29 and 34 are omitted from the standard Schmidt formulation.[23]

Hierarchic Clustering

The reduced order methods presented in the previous sections provide viable solutions to the fleet navigation problem for medium-sized fleets. For larger fleets, however, the *synchronization requirements* will

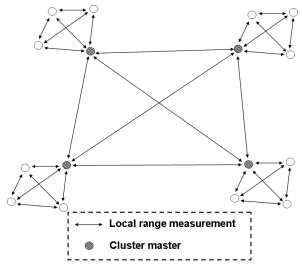

Fig. 4: Hierarchic Clustering Topology.

make the iteration techniques very difficult and an alternative architecture will be required. One approach to mitigate this scaling problem, is to employ a hierarchic architecture, which performs the detailed estimation for smaller groups of vehicles and then assimilates partial results at a higher level. Fig. 4 illustrates one strategy for setting up a hierarchic architecture for spacecraft employing local ranging. The fleet is split up into smaller clusters that perform their ranging and navigation independently with the exception of one vehicle in each cluster termed the "cluster master". To link the estimates of each cluster to one another, each cluster master joins together to form a "super-cluster", so the hierarchy looks essentially identical at each level. The key benefit of this approach is that the clusters and super-clusters do not need to be tightly synchronized, or even run at the same update rate.

The type of filter for the cluster estimators is chosen based on the cluster sizes, available communication bandwidth, CPU loading and required accuracy. Depending upon the number of layers in the architecture (only two are shown in Fig. 4), determining a vehicle's position with respect to the fleet center may involve vector additions of several different cluster solutions. Since each vector addition step introduces another source of error (due to summing two quantities, each with an uncertainty), from a performance perspective, larger clusters are more desirable than smaller ones. Assuming a two-layer hierarchy, and each cluster running identical estimators, the predicted growth of estimation error variance is shown in Table 1.

Cluster sizing and selection could be done based on several different criteria, including geographic separation, common GPS visibility or even existing communication connectivity from science experiments. If the fleet communication architecture permits, the clusters could be dynamic. From an estimation standpoint, the best clustering approach would be to form clusters of vehicles that are ranging with each other; however, inter-cluster ranging could be permitted

Table 1: Error scaling through hierarchy with σ^2 as basic estimator position error variance. Errors accumulate across the levels in the hierarchy.

Estimate	**Error Variance**
Slave to Local Master	σ^2
Slave to Local Slave	$2\sigma^2$
Local Master to Local Master	σ^2
Local Master to Fleet Center	σ^2
Slave to Fleet Center	$2\sigma^2$
Slave to Remote Master	$3\sigma^2$
Slave to Remote Slave	$4\sigma^2$

provided enough state information is exchanged to formulate the measurements. For example, if a vehicle in one cluster wants to range with another vehicle in different cluster, information must be traded regarding the other cluster's global position (*i.e.,* must communicate with the cluster masters). However, if inter-cluster ranging is limited to only cluster masters, then the only extra information required is known by the local cluster master which substantially reduces the communication requirements. Performing inter-cluster ranging is a good way to reduce the error growth illustrated in Table 1 since the error variances associated with positions between ranging vehicles is at most $2\sigma^2$.

Hierarchic clusters permit much greater flexibility in terms of the estimation algorithms – any of the previously algorithms could be used at any level of the hierarchy. Choosing the most appropriate algorithms to populate the various levels depends on the required estimation accuracy, the available computation, and the communication bandwidth.

Analysis

The following sections present performance analysis results for the filters discussed in the previous sections. This includes an analysis of the steady state filter covariances. To verify the predicted errors, simulations are also done for each method to obtain representative accuracy estimates. Finally, results are presented from tests run on the FFIT testbed to analyze the overall performance of each method from a data-flow point of view.

Covariance Analysis

The focus of this study was to observe the structure of the filter covariance matrix in each filter. Constant probability contours are a useful visualization technique for studying covariance matrices. In the

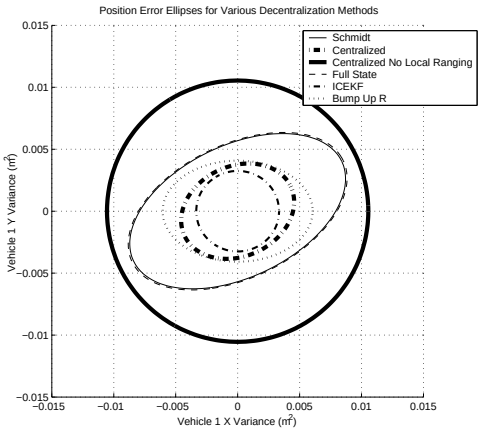

Fig. 5: Error ellipses for various decentralization methods (2D Case)

2×2 case, a covariance matrix can be represented as a rotated ellipse in the Cartesian plane.

For this study, each vehicle receives an x and y measurement for their own position as well as the range between the vehicles, r_{12}. The algebraic Riccati equation is solved for each filter to obtain the steady state covariance matrix. Fig. 5 shows the error ellipses for vehicle 1 for each type of centralized and decentralized filter described above (since the Kalman and Information filters provide identical covariances, only one ellipse is shown for both). This type of analysis provides insight into the degree to which the measurements are being used. For example, a narrow ellipse aligned with the x-axis would indicate that the filter had been able to make very good use of a measurement in the y-direction (resulting in better confidence and hence a narrower ellipse, in the y-direction).

The ellipse representing the case with no local ranging is a circle with a radius larger than any of the semi-major axes of the other ellipses. The circular shape is expected due to the equal measurement accuracy in the linear positional measurements of x and y. The ellipse for the centralized filter represents the *best case* possible since the centralized method captures all vehicle and state correlation in a single, unified filter. Both the *Bump Up R* method and the *ICEKF* method appear to have covariance matrices that are relatively close to that of the centralized method, suggesting near-centralized performance, but this is misleading because the filter covariance for the *Bump Up R* and *ICEKF* methods do not provide good figures of merit, as outlined in the following.

In a full order, centralized Kalman filter, the measurement equation is

$$z = CX \tag{35}$$

where X denotes the entire fleet state and z denotes the full complement of measurements available in the fleet. This equation can be expanded into terms that correspond to the local states x and the states of the other vehicles y

$$z = CX = Hx + Jy \tag{36}$$

Note that this partitioning is identical to that used to derive the SKF. The ICEKF method is a reduced-order estimator so it only has the local state available to use in the measurement equation. So its measurement equation is given by the approximation

$$z_i \approx H_i x \tag{37}$$

where z_i is the vector of measurements available on a local vehicle. The term Jy is completely omitted from the ICEKF equations; however, the regular Kalman filter equations derived assuming an optimal K with correct values of R, Q, Φ and H are still used. This means that the filter covariance cannot be trusted as the true error covariance and thus is not a good indicator of the filter performance. Furthermore, the fact that the ICEKF covariance ellipse is smaller than the centralized ellipse (*i.e.,* the *true* covariance[‡]) indicates that the ICEKF method is trusting the measurements too much. Similarly, the *Bump Up R* covariance cannot be trusted either since here, not only is the H incorrect, but the R is incorrect as well.

The filter covariance of the full state decentralized estimator, however, can be trusted since it uses the full H matrix and the correct R value. Notice that in one direction, it provides only a minimal improvement over the *No Local Ranging* case. Also, since this filter attempts to estimate many more states than are observable, the covariance corresponding to the states not pertaining to the local vehicle will tend to grow without bound (depending on the dynamics, Φ). This could lead to an ill-conditioned P matrix and cause numerical problems over time if not adequately accounted for.

The SKF uses a sub-optimal K but the filter co-

[‡]Actually, since these are Extended Kalman Filters, the *true* error covariance is not accessible. However, the covariance error introduced in using an EKF formulation over that introduced by using an incomplete measurement matrix can be shown to be small if the measurement linearization works sufficiently well.

variance *can* be trusted. Recall that the SKF derivation starts with the general update equation (Eq. 38) for any (optimal or sub-optimal) K.

$$P_k^+ = (I - K_k H_k) P_k^- (I - K_k H_k)^T + K_k R_k K_k^T \quad (38)$$

Using this equation, the true covariance is recovered and hence, the SKF covariance can be trusted. The ellipse for the SKF method appears only slightly better than the full state case. This is somewhat surprising based on the correction that the SKF attempts, but the SKF does give a reduced computational load (as seen in the Analysis Section). Also the SKF does not suffer from the ill-conditioned P matrices that are typically experienced with the full state filters.

These results indicate that filters such as the centralized and SKF methods should perform well since the filter's conception of the error covariance is accurate. Methods such as ICEKF are not expected to perform well since they over-emphasize the quality of the measurements, as evidenced by small covariances. Furthermore, being able to trust the filter's covariance provides valuable sensor integrity information which could be used in fault detection routines (*e.g.,* more ranging partners may be required if the filter covariance becomes too high).

Simulation Results

To demonstrate the relative effectiveness of the various filters presented in this paper, simulations of a 7-vehicle fleet consisting of one "master" vehicle and 6 "slave" vehicles were conducted. The goal of each slave vehicle is to estimate its relative state with respect to the master vehicle. Since only relative states are of interest for this particular simulation, the absolute state of the master is taken arbitrarily to be at the origin. Each slave follows different circles of constant radii and angular speeds around the master to provide good fleet geometry (PDOP) for utilizing the local ranging measurements most effectively.

For estimation purposes, there are two types of measurements available. The first type are coarse measurements of the relative position and velocity vectors with respect to the master. These measurements are intended to represent the results from some basic navigation system (*i.e.,* NAVSTAR-only GPS). The second type are fine local ranging measurements. Ranges are determined by measuring the time-of-flight of a signal, requiring knowledge of the differential clock offsets of each vehicle pair. These clock offsets must be incorporated into each vehicles' state vector for estimation (see Park[20] for more details). Range rates are determined from the Doppler measurements. The coarse measurements are assumed to be accurate to 3 cm and 7 mm/s, and the fine measurements are a factor of 3 more accurate.

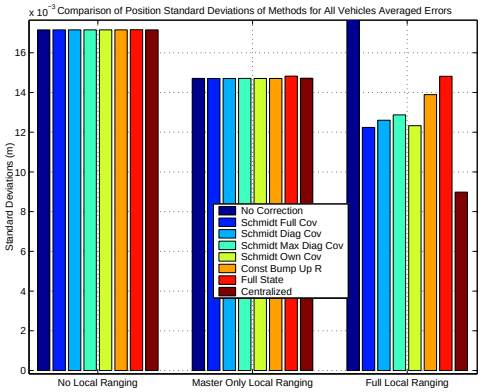

Fig. 6: Bar chart illustrating position estimation error for various decentralization methods (ICEKF bar is beyond top of chart). Errors reported as root-sum-squared, averaged over multiple runs.

The simulation has three measurement phases. In the first phase, the only measurements available to each vehicle are the coarse position and velocity measures. The second phase adds local ranging measurements, but only between the master and slaves (not between slaves). The final phase includes the full complement of coarse and fine measurements including local ranging between slaves. Fig. 6 shows the results of the simulation during each phase. The methods tested include (in order from left to right) ICEKF, 4 types of SKF, Bump-Up-R, Full State and Centralized. Four SKF methods were tested:

- **Schmidt Full Cov** - The full P_{xx} for each vehicle is exchanged at each time-step to be used in each other vehicles' P_{yy}.
- **Schmidt Diag Cov** - This approach saves some inter-spacecraft communication by transmitting only the diagonal of each vehicle's P_{xx} to form the P_{yy} matrix.
- **Schmidt Max Diag Cov** - This method is similar to the previous method, but only the maximum diagonal element of each vehicles' P_{xx} is transmitted. To form the P_{yy} matrix, each vehicle multiplies this maximum diagonal element by the identity matrix.
- **Schmidt Own Cov** - In this final SKF method, each vehicle assumes that its own covariance matrix, P_{xx}, is similar to the covariance matrices of all other vehicles in the fleet. So, nothing is transmitted and each vehicle simply uses its own P_{xx} to form P_{yy}.

Fig. 6 compares the performance of each decentralized method to that of the centralized method. The three clusters of bars in the charts represent the three different phases in the estimation process. All decentralization methods produce the same accuracy as the centralized for the first two phases corresponding to no local ranging and master-only local ranging. This result is expected since the H matrix for both of these phases is block diagonal and thus the problem cleanly decentralizes. It is also clear from the Fig. 6 the effect that the local ranging has on the best achievable accuracy.

For the third phase, the H matrix is no longer block diagonal and the differences between the decentralization methods become evident. In particular, the poor results of the ICEKF with no correction (and only 1 iteration) are readily apparent. Increasing the number of iterations can improve this result, but the fundamental problem is that the approach uses the measurements too aggressively. Of the methods tested, the SKF method with full covariance transmission performs the best, second only to the centralized method. While the best SKF case is worse than the centralized result, Fig. 6 shows that it outperforms the "Bump up R" approach. Note that the more covariance information traded, the better the achieved performance. Another interesting observation is how well the "Schmidt Own Cov" case performs. Due to similar geometries, it is likely that the covariances look similar for each vehicle, providing accuracies almost as high as the "Schmidt Full Cov" case.

Data Flow Validation

The FFIT testbed[15] enables data to be collected on the communication and computational demands for each algorithm. The relative complexity of the algorithms can also be studied from the time required to generate a useful solution. The FFIT testbed currently contains 5 PCs in total - 4 spacecraft PCs and 1 simulation engine, so the simulations are limited to a fleet of 4 vehicles. While this is not as large as desired, the results provide a baseline that can be extrapolated for larger fleets (*e.g.,* 16–32 vehicles). The results that follow were taken directly from the FFIT testbed using a situation similar to the one in the previous section with some simplifications (*i.e.* 2-D instead of 3-D and all clock offsets were assumed known). The results were taken in real-time at a time-step of 5 seconds, with each vehicle transmitting its data at a rate of 9600 baud (typical data rate for a Nanosat missions[22]). Fig. 7 shows the trade between estimation accuracy and

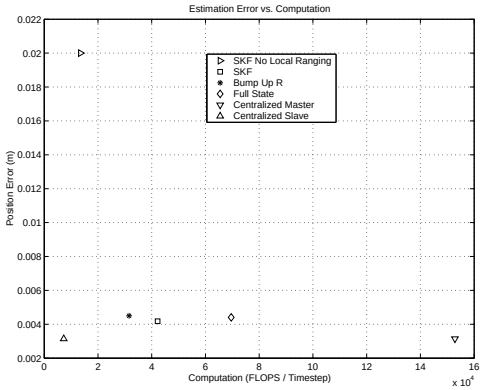

Fig. 7: Estimation error as a function of computation per vehicle for various decentralization methods.

computation. The centralized method provides the lowest error, but this plot clearly shows the computational price that is paid for this high accuracy. The master vehicle in the centralized estimation scheme ends up more than 10 times more computationally loaded than the slaves. In a mission where science objectives need to be carried out along with relative navigation, it may not be desirable to have one vehicle so heavily loaded while the others are doing virtually nothing.

The reduced order methods (*i.e.,* all SKF methods and bump-up-R) provide a good trade of accuracy vs. computation plot (especially compared to the "no local ranging" data point). Since each vehicle only estimates its local state, the computation is well distributed and as a result, the computation level for each vehicle is only slightly more than a centralized slave. The results in Fig. 7 show that the full-state method has a higher computational load than the other reduced order methods (*e.g.,* full SKF method), but this extra computation provides no additional performance benefit.

Overall, the other reduced-order methods require relatively small amounts of inter-spacecraft communication as evidenced in Fig. 8. The method with the most communication is the SKF due to the transmission of portions of covariance matrices. However, given the accuracy benefit of the SKF, it may be desirable to use more communication bandwidth in exchange for the increased estimator accuracy. The centralized method does very well in this comparison because the fleet only has 4 vehicles and only the measurement vectors and solutions for the two-dimensional problem are transmitted. The issues when scaling to larger fleets are explored later in this section.

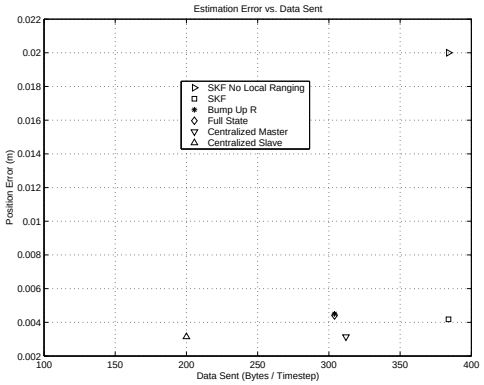

Fig. 8: Estimation error as a function of transmitted data amounts per vehicle for various decentralization methods.

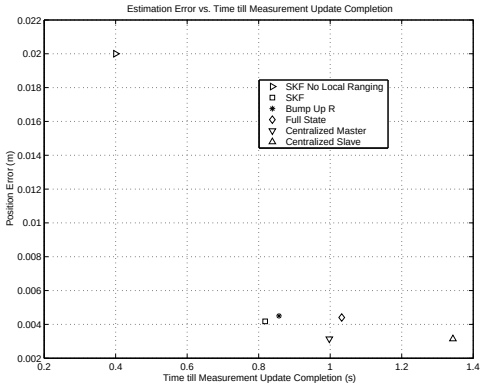

Fig. 9: Estimation error as a function of solution time per vehicle for various decentralization methods.

Fig. 9 shows the estimation error vs. solution time trade. This plot makes explicit the degree of complexity of each method. Points further to the right indicate that a great deal of computation, communication and/or fleet synchronization is required, thus making the method more complicated to execute overall. The reduced order methods appear to be the clear winners in this trade study for $N = 4$. They minimize computation as well as communication and still perform quite well from an estimation accuracy standpoint. Furthermore, due to the cascade nature of the measurement updates, each vehicle need only be synchronized with the vehicle before it in the fleet. In the centralized case, all vehicles must synchronize with the master prior to receiving a measurement update, resulting in longer solution times.

The results presented above for a 4-vehicle fleet provided insights on the relative merits of the various estimation architectures. The reduced-order methods exhibit near optimal estimation accuracy while keeping communication and computation to a minimum, proving that decentralized estimation is possible even with correlated states. The full state methods were shown to not provide any performance increase over the reduced-order methods in terms of computation or estimation accuracy. From a communication and computation perspective, for these simplified simulations, the centralized filter appears to be a viable option for fleets of only 4 spacecraft, provided the imbalance in computational load is acceptable. However, this analysis has not included any notion of fleet robustness or required fleet connectivity, which are the commonly cited disadvantages of centralized estimators.[17,19]

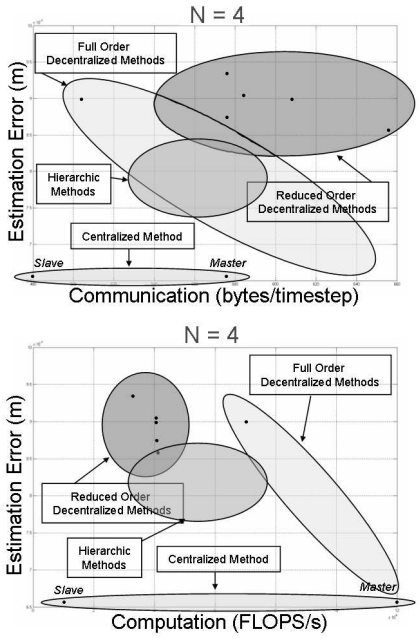

Fig. 10: Trends for $N = 4$.

To gain some insight into how these algorithms scale to a larger fleet, Figs. 10 and 11 use the $N = 4$ case as a baseline and illustrate graphically how the performance of the various estimation architectures changes when the fleet size increases past 4 vehicles. The scaling presented in these figures comes from an algebraic analysis of the required matrix calculations.[25] Fig. 10 illustrates the approximate regions on the communication and computation plots where the various architectures lie. The "Full State" area stretches down to the bottom right on both plots because it includes the information filter as well. Fig. 11 shows how these regions evolve when for N larger than 4 (*i.e.* 16 or 32).

In the communication plot, the scaling for the full state and centralized methods grows with N^2 due to the need to send information matrices in the information filter case and full measurement vectors

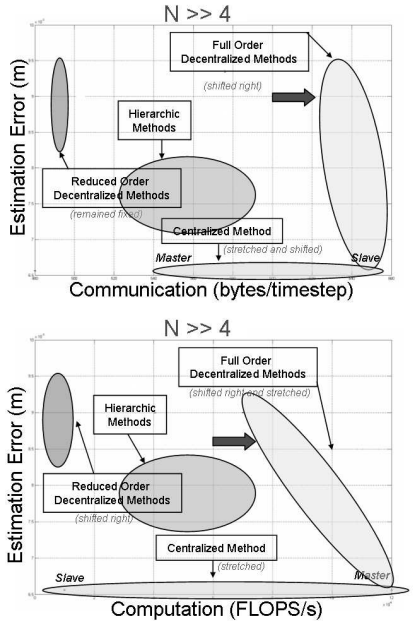

Fig. 11: Trends for $N \gg 4$.

in the centralized Kalman filter architecture. For the reduced-order methods, there is *no* communication scaling with N since all transmitted data is only a function of the local vehicles' states. In the computation plot, the scaling for the full order methods grows with N^3, while the scaling is N^6 for the centralized methods. The reduced-order methods however, experience no scaling in computation.

The scaling results presented above indicate that the reduced-order methods perform the best in larger fleets due to their limited state size. The communication and computation requirements for the centralized and full state methods grow extremely rapidly for larger fleets, quickly rendering those methods infeasible. However, all presented methods require some degree of fleet synchronization, which may limit the implementation for large N.

Due to the physical limitations of the FFIT testbed, it is not possible to simulate the hierarchic method described in this paper. However, since the hierarchic methods are simply other estimators running on various clusters, the above results can be extrapolated to the hierarchic architectures. If one assumes that the same estimator is run at each level of the hierarchy, then the communication and computational load for each cluster master would be exactly twice that of the cluster slaves. The slaves' communication and computational load scale exactly the same as predicted in Figs. 10 and 11. Thus, the hierarchic methods permit scaling mitigation by simply reducing the size of each cluster.

An important drawback to hierarchic clustering is the summation of cluster error variances shown in Table 1. Of course, to reduce the impact of these errors, larger clusters can be used. Using reduced-order estimators such as the SKF permit large cluster sizes due to their appealing scaling characteristics. Thus, a viable navigation option for almost any large fleet would be a reduced order estimator implemented in a hierarchic cluster architecture.

Conclusions

Estimation for formation flying missions at MEO and beyond presents a challenging problem due to the nonlinear ranging measurements that will be required. These local ranging measurements strongly couple the states of each vehicle and complicate the decentralization of the estimation algorithms by requiring cascaded fleet iterations during the measurement update. This paper analyzed several estimation architectures and compared various algorithms using 2D and 3D simulations that facilitated detailed studies of the effects of the nonlinearity in the ranging measurements and the correlation between the vehicle state estimates. Results from these simulations showed that the decentralized reduced-order estimators provide a good balance between communication, computation and performance when compared to centralized and full order methods, and thus could be a feasible relative navigation approach for future missions. The extrapolation of these results to larger fleets strongly indicated that the centralized (and decentralized full-order) filters will have prohibitively high communication and computational requirements. However, the reduced-order estimators presented, of which the Schmidt-Kalman filter was the best, exhibit no such growth in the communication or computation demands. While the reduced-order decentralized approaches reduce and distribute the computation more equitably, they are fundamentally limited by the degree of synchronization required within the fleet to calculate the state estimates. The hierarchic architectures discussed address this limitation by splitting the fleet into subteams that can function asynchronously. Thus, the results of this study show that a viable estimation approach for large fleets using augmented measurements would be comprised of reduced-order estimators implemented within a hierarchic architecture.

Acknowledgments

Research funded under NASA grants NAG5-10440 and NAG3-2839.

References

[1] F. H. Bauer, K. Hartman, E. G. Lightsey, "Spaceborne GPS: Current Status and Future Visions," proceedings of the *ION-GPS Conference*, Sept. 1998, pp. 1493–1508.

[2] F. H. Bauer, K. Hartman, J. P. How, J. Bristow, D. Weidow, and F. Busse, "Enabling Spacecraft Formation Flying through Spaceborne GPS and Enhanced Automation Technologies," *ION-GPS Conference*, Sept. 1999, pp. 369–384.

[3] J. Leitner, F. Bauer, D. Folta, R. Carpenter, M. Moreau and J. How, "Formation Flight in Space," *GPS World,* Feb. 2002, pp. 22–31.

[4] R.W. Beard, J. Lawton and F.Y. Hadaegh, "A Coordination Architecture for Spacecraft Formation Control," *IEEE Transactions on Control Systems Technology*, July, 2000.

[5] F. D. Busse, *Precise Formation-State Estimation in Low Earth Orbit using Carrier Differential GPS*. Ph.D. thesis, Stanford University, Dept. of Aero/Astro, Nov. 2002.

[6] F. Busse and J. How, "Demonstration of Adaptive Extended Kalman Filter for Low Earth Orbit Formation Estimation Using CDGPS", *ION-GPS*, Sept. 2002.

[7] F. Busse and J.P. How, "Real-time Experimental Demonstration of Precise Decentralized Relative Navigation for Formation Flying Spacecraft", *AIAA GNC*, August 2002.

[8] T. Corazzini, *Onboard Pseudolite Augmentation for Spacecraft Formation Flying*. Ph.D. Dissertation, Stanford University, Dept. of Aeronautics and Astronautics, Aug. 2000.

[9] C.-W. Park, J. How, and L. Capots, "Sensing Technologies for Formation Flying Spacecraft in LEO Using Inter-Spacecraft Communications System," the *Navigation* J. of the Institute of Navigation, Vol. 49, No. 1, Spring 2002, pp. 45–60.

[10] J. R. Carpenter, C. Gramling *et al*, "Relative Navigation of Formation-Flying Satellites," from the proceedings of the *International Symposium on Formation Flying*, Toulouse France, October 2002.

[11] P. Stadter, R. Heins, et al., "Enabling Distributed Spacecraft Systems with the Crosslink Transceiver," *AIAA Space Conference and Exposition*, Aug. 2001. Paper 2001-4670.

[12] G. Purcell, D.Kuang, S. Lichten, S.C. Wu and L. Young, "Autonomous Formation Flyer (AFF) Sensor Technology Development," TMO Progress Report 42-134, August 1998.

[13] J. R. Carpenter, "Decentralized control of satellite formations," *Int. J. Robust Nonlinear Control* 2002; 12:141-161.

[14] C. Park and J. P. How, "Precise Relative Navigation using Augmented CDGPS," *ION-GPS Conference*, Sept. 2001.

[15] P. Ferguson, T. Yang, M. Tillerson and J. How, "New Formation Flying Testbed for Analyzing Distributed Estimation and Control Architectures." Presented at the *AIAA GNC*, Monterey, CA, August 2002.

[16] H.F. Durrant-Whyte. "Sensor models and multi-sensor integration," *Int. J. Robotics Research*, 7(6):97-113, 1998.

[17] A. G. O. Mutambara *Decentralized Estimation and Control for Multisensor Systems*, CRC Press LLC, 1998.

[18] J.L. Speyer, "Communication and transmission requirements for a decentralized linear-quadratic-gaussian control problem," *IEEE Trans. Automatic Control*, 24(2):266-269, 1979.

[19] J. R. Carpenter, "Partially Decentralized Control Architectures for Satellite Formations,", Presented at the *AIAA GNC Conf.*, Monterey, CA, Aug. 2002.

[20] C. W. Park, *Precise Relative Navigation using Augmented CDGPS*. Ph.D. thesis, Stanford University, Dept. of Mech. Eng., June 2001.

[21] S. F. Schmidt, "Application of State-Space Methods to Navigation Problems," in C. T. Leondes (ed.), *Advances in Control Systems*, Vol. 3, New York: Acad. Press, 1966.

[22] P. Ferguson, F. Busse, *et al*, "Formation Flying Experiments on the Orion-Emerald Mission", *AIAA Space 2001 Conference and Exposition*, Albuquerque, NM, Aug. 2001.

[23] P. Ferguson, *Distributed Estimation and Control Technologies for formation Flying Spacecraft*. S.M. thesis, Massachusetts Institute of Technology, Dept. of Aero/Astro, Jan. 2003.

[24] R. Brown, P. Hwang, *Introduction to Random Signals and Applied Kalman Filtering*. Third Ed., pp. 366-7, John Wiley & Sons, 1997.

[25] M. S. Grewal and A. P. Andrews, *Kalman Filtering Theory and Practice*, pp. 208-210, Prentice-Hall Inc., 1993.