

Federated Kalman Filter Simulation Results

NEAL A. CARLSON

Integrity Systems, Inc., Belmont, Massachusetts

MICHAEL P. BERARDUCCI

Air Force Wright Laboratory, WPAFB, Ohio

Received June 1993

Revised March 1994

ABSTRACT

This paper describes federated filter applications to integrated, fault-tolerant navigation systems, with emphasis on real-time implementation issues and numerical simulation results. The federated filter is a near-optimal estimator for decentralized, multisensor data fusion. Its partitioned estimation architecture is based on theoretically sound information-sharing principles. It consists of one or more sensor-dedicated local filters, generally operating in parallel, plus a master combining filter. The master filter periodically combines (fuses) the local filter solutions to form the best total solution. Fusion generally occurs at a reduced rate, relative to the local measurement rates. The method can provide significant improvements in fault tolerance, data throughput, and system modularity. Numerical simulation results are presented for an example multisensor navigation system. These results demonstrate federated filter performance characteristics in terms of estimation accuracy, fault tolerance, and computation speed.

INTRODUCTION

Integrated multisensor navigation systems have the potential to provide high levels of accuracy and fault tolerance. The presence of multiple data sources provides functional redundancy, as well as greater observability of the desired navigation states. However, that potential has not always been fully realizable through the application of standard (centralized) Kalman filtering techniques for multisensor data fusion. Centralized filters can result in severe computation loads when processing data from multiple sensors. Worse, when used in two-stage (cascaded) filter architectures, standard Kalman filters can exhibit poor accuracy and unpredictable behavior under some conditions.

During the past 15 years, the development of decentralized (or distributed) Kalman filtering methods has received increasing attention. Parallel processing technology, emphasis on fault-tolerant system design, and availability of multiple specialized sensors strongly motivate the development of such methods. Potential applications include multisensor navigation systems, tracking systems, and other data fusion systems.

Early contributions to optimal decentralized filter theory [1-5] provided useful insights about different approaches to partitioned estimation. Building

on this theoretical base, reference [6] proposed a decentralized filtering structure with attractive fault-tolerant features. Reference [7] developed an optimal and practical decentralized filtering method for orbit estimation purposes. References [8] and [9] made subsequent extensions of that approach for other applications.

More recently, the federated filter method based on rigorous information-sharing principles was developed [10–12]. This method provides globally optimal or near-optimal estimation accuracy with a high degree of fault tolerance, and is practical for real-time distributed navigation systems applications [13]. The federated filter structure employs sensor-dedicated local filters (LFs), and a master filter (MF) to combine or fuse the LF outputs. The method allows several different information-sharing strategies, or modes, for different applications. Some modes require simultaneous updates from the LFs to the MF, while others permit independent updates from each LF.

Recently, the net information approach to decentralized estimation was developed [14–15]. This globally optimal/near-optimal method is based on information-sharing principles similar to those of the federated method. Each local element employs two LFs (one is locally optimal) and a differencing algorithm to determine the new information gained over each period. Local elements pass their new information to the MF independently, for data fusion.

This paper focuses on the performance of the federated filtering technique applied to integrated, fault-tolerant navigation systems. The remaining sections describe the distributed estimation problem, limitations of standard Kalman filters, federated filter basics, federated configuration options and features, numerical simulation results, and conclusions.

DISTRIBUTED ESTIMATION PROBLEM

The federated filter is a partitioned estimation method. It employs a two-stage (cascaded) data processing architecture, in which the outputs of sensor-related LFs are subsequently combined by a larger MF, as illustrated in Figure 1.

As indicated, each LF is dedicated to a separate sensor subsystem, and also uses data from a common reference system, generally an inertial navigation system (INS).

The federated filter technique comprises a weighted least-squares solution to the following linear (or linearized) estimation problem. First, consider a system state vector \mathbf{x} that propagates from time t' to $t = t' + \Delta t$ according to the following dynamic model:

$$\mathbf{x} = \Phi \mathbf{x}' + \mathbf{G} \mathbf{u} \quad (1)$$

Here, Φ is the state transition matrix between time points t' and t , \mathbf{G} is the process noise distribution matrix, and \mathbf{u} is the additive uncertainty vector due to white process noise acting over the timestep Δt . The errors \mathbf{e}_0 in the initial state estimates $\hat{\mathbf{x}}_0$ have covariance \mathbf{P}_0 . The sequential process noise values \mathbf{u}_j at times t_j have covariances \mathbf{Q}_j , and are uncorrelated with \mathbf{e}_0 and with each other.




Fig. 1—General Federated Filter Structure

Our system also has access to external measurements \hat{z}_i from $i = 1 \dots n$ separate local sensor subsystems. The discrete measurements from sensor # i at time t , are related to the true state x_t by a linear (or linearized) relationship:

$$\hat{z}_i = H_i^T x_t + v_i \quad (2)$$

Here, H_i^T is the sensor # i measurement observation matrix (often defined without the transpose), and v_i is the additive, random measurement error. The sequential error values v_i have covariances R_i , and are uncorrelated at successive times t . In addition, measurement errors from different sensors i and m are statistically independent. Disjoint sensor data sets permit the total estimation problem to be divided into sensor-related partitions with independent measurement processes.

LIMITATIONS OF STANDARD KALMAN FILTERS

The primary limitations of standard (centralized) Kalman filtering methods when applied to multisensor navigation systems and/or systems with embedded local filters are (1) heavy computation loads, (2) poor fault-tolerance, and (3) inability to correctly process prefiltered data in a cascaded (two-stage) filter structure.

The first limitation of a centralized Kalman filter (CF) is fairly obvious. In Figure 1, the presence of several sensors generally implies a relatively large number of filter states n , since each sensor typically introduces one to five measurement bias states. For a single CF, the per-cycle computation load grows roughly in proportion to $n^3 + \sum m \cdot n^2$, where $\sum m$ is the total number of

measurements across all the sensors. The problem is especially severe when one or more sensors require high-rate measurement processing. (However, computation load is a less critical concern with modern, high-speed processors and algorithms that take advantage of sparseness in the Φ , Q , and H matrices.)

The second limitation of a CF relates to fault tolerance. Like any optimal filter, the CF attempts to make its data inputs agree in a weighted least-squares sense, thereby suppressing their differences. An undetected failure in one sensor gets distributed into all of the navigation state and sensor bias estimates, so that they all tend toward agreement. Thus, while measurement residual tests can readily detect sudden *hard* failures, they may completely miss gradual *soft* failures. If the CF does incorporate faulty data from any sensor, its full solution becomes corrupted and must be reinitialized.

The third limitation of a CF relates to cascaded filter processing, and can best be illustrated by means of an example. Figure 2 shows the major components of a cascaded filter designed for an integrated navigation system composed of an INS, a GPS Phase IIIA receiver/navigator, and a radar subsystem. GPS receiver measurements and INS outputs are processed locally by an embedded GPS/inertial Kalman filter. Periodically, position and velocity outputs from the GPS LF are incorporated as discrete "measurements" by a Kalman MF in the central computer.

One particular aspect of this ad hoc cascaded filter design can lead to accuracy and/or stability problems, given a Kalman MF: errors in the position and velocity outputs from the GPS LF are sequentially correlated, not uncorrelated as a Kalman filter requires. Sequential correlations in the GPS filter outputs imply that each output contains some new and some old information. Treating each GPS filter output as entirely new information causes the MF to become overoptimistic regarding its own accuracy (i.e., its covariance gets too small). The resulting problem is especially obvious if we suppose that GPS receiver

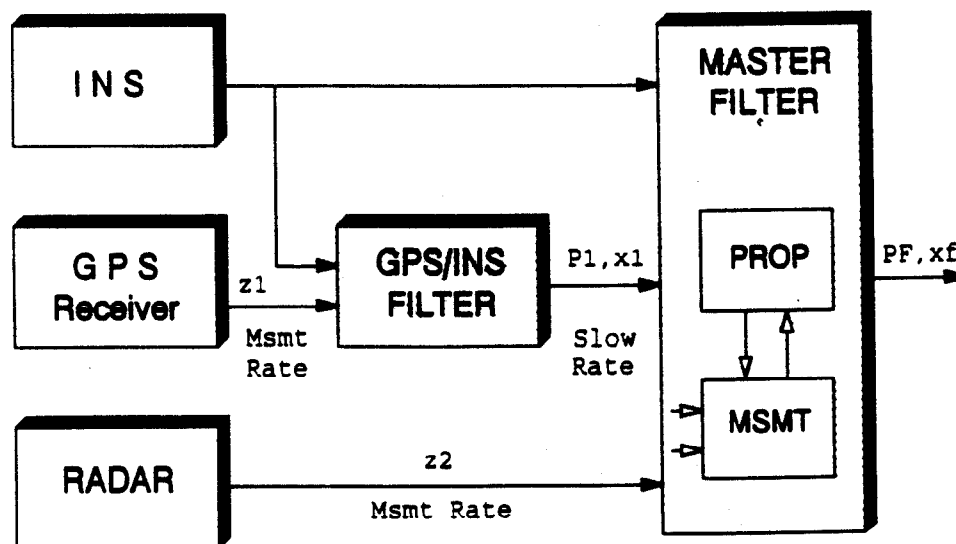


Fig. 2—Ad Hoc GPS/Radar Filter

measurements become unavailable for several minutes. The MF will continue to incorporate successive GPS filter outputs as fresh "measurements" and to reduce its covariance accordingly, even though those outputs contain absolutely no new information.

A common, ad hoc fix for these correlation problems is to limit the GPS incorporation rate in the Kalman MF to once every 10 or 20 s. Over these intervals, the GPS Phase IIIA filter output errors typically become decorrelated, so that they appear sequentially random to the MF. (Rapid decorrelation is caused by large process-noise terms in the GPS filter.) This ad hoc cascaded filter design can yield satisfactory estimation performance for many applications. However, questions remain as to whether the same approach will work with different GPS filter tuning parameters or with other sensor subsystems.

FEDERATED FILTER BASICS

The new federated filter method avoids the theoretical and practical difficulties described in the previous section by means of a simple yet effective information-sharing methodology. We use the term *information* to denote the filter solution, i.e., the filter state vector and covariance matrix (or equivalently its inverse, commonly called the information matrix). The basic concept of the information-sharing approach implemented by the federated filter is this: (1) divide the total system information among several component (local) filters; (2) perform local time propagation and measurement processing, adding local sensor information; and (3) recombine the updated local information into a new total sum.

The remainder of this section briefly reviews how the federated filter applies information-sharing principles in its use of the n LFs and one MF in Figure 1. The presentation is heuristic, to emphasize the general concept. References [10, 11, 16] provide a rigorous theoretical derivation of the federated filter.

First, let the full, CF solution be represented by the covariance matrix $\mathbf{P}\mathbf{F}$ and state vector $\hat{\mathbf{x}}\mathbf{f}$; the i th LF solution by $\mathbf{P}\mathbf{I}$ and $\hat{\mathbf{x}}\mathbf{i}$; and the MF solution by $\mathbf{P}\mathbf{M}$ and $\hat{\mathbf{x}}\mathbf{m}$. We will use index $i = 1 \dots n$ for the LFs alone, and $k = 1 \dots n, m$ for the MF plus LFs, where $k = m$ is the MF. Now, if the LF and MF solutions are statistically independent, they can be optimally combined by the following additive information algorithm, where the inverse covariance matrix \mathbf{P}^{-1} is known as the information matrix:

$$\mathbf{P}\mathbf{F}^{-1} = \mathbf{P}\mathbf{M}^{-1} + \mathbf{P}\mathbf{I}^{-1} + \dots + \mathbf{P}\mathbf{N}^{-1} \quad (3)$$

$$\mathbf{P}\mathbf{F}^{-1}\hat{\mathbf{x}}\mathbf{f} = \mathbf{P}\mathbf{M}^{-1}\hat{\mathbf{x}}\mathbf{m} + \mathbf{P}\mathbf{I}^{-1}\hat{\mathbf{x}}\mathbf{1} + \dots + \mathbf{P}\mathbf{N}^{-1}\hat{\mathbf{x}}\mathbf{n} \quad (4)$$

The key to the new federated filtering method is to construct individual LF and MF solutions so they can be combined or recombined at any time by the above simple algorithm. In particular, the construction avoids the need to maintain local/local or local/master cross-covariances. The procedure for doing so is the essence of the federated filter method.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.