R.C. Raup, R.A. Ford, G.R. Krumpholz, M.G. Czerwinski, and T.E. Clark

The Best Approximation of Radar Signal Amplitude and Delay

The estimation of receiver signal amplitude and delay, which can be converted to target cross section and range, is one of the fundamental functions of signal processing algorithms in a narrowband radar. The problem of tracking high-earth-orbit (HEO) satellites with ground-based radars requires a generalization of simple filter-bank signal processing architectures traditionally used to estimate signal amplitude and delay. A solution of the mathematical best-approximation problem leads to a new signal processing architecture that efficiently estimates signal amplitude and delay in all of the generality necessary to address the HEO satellite tracking problem.

Narrowband radars capable of tracking high-earth-orbit (HEO) satellites are an important resource for both present and future space missions. To track satellites, the United States maintains a worldwide network of ground-based sensors that provide data for civilian and government users. At any given time, one or more ground-based radars are engaged in HEO satellite tracking. This routine tracking becomes more important as the number of man-made objects in space increases. Ground-based sensors regularly track more than 5000 space objects, and hundreds more are added each year.

The orbits and locations of active satellite payloads are well known because of telemetry or transponder-based tracking. The majority of space objects, however, are inactive. If these inactive satellites are not periodically tracked and their orbit and position carefully measured, their location would soon become unknown. Most of the tracking is performed by ground-based optical and narrowband radar sensors. Knowing the locations of these objects is important to reduce the risks of collision in space, particularly with manned spacecraft, and to maintain national security.

Satellites can be divided into two types: low earth orbit and high earth orbit, depending on their average altitude and orbital period. Each type presents a unique challenge to the designers of ground-based radars. The designer must provide algorithms to estimate the radar signal amplitude and delay, which are then used to determine target cross section and range. This article describes a solution to the estimation problem, developed to support the unique requirements of HEO satellite tracking. The solution, which is derived from a mathematical best-approximation viewpoint, leads to a generalization of traditional receiver designs. The solution can be efficiently implemented with modern computational methods.

The High-Earth-Orbit Satellite Tracking Problem

The HEO target tracking problem places constraints on the design of signal processing systems in ground-based radars. These constraints are imposed by low SNRs, and significant pulse-to-pulse amplitude variations in the received signal. The low SNRs occur when small targets are tracked at long ranges. The amplitude variations occur when the effects of change in target aspect angle combine with the complex scattering patterns of the satellites to cause fluctuation in apparent target cross section. The histograms shown in Fig. 1 illustrate these two characteristics of HEO satellite tracking.

The histogram in Fig. 1(a) characterizes the average single-pulse SNR of 722 targets tracked by the Millstone Hill L-band radar in Westford, Mass. The fraction of the total number of targets per dB of SNR is plotted versus the average single-pulse SNR. The radar system designer typically requires at least 10 dB of SNR to detect the presence of the target. Because more than half of the area under the histogram curve is associated with SNRs below 10 db, most targets cannot be detected by the sensor on the basis of a single radar pulse. Hence multiple-pulse detection and estimation signal processing play an important role in the design of a HEO satellite surveillance radar.

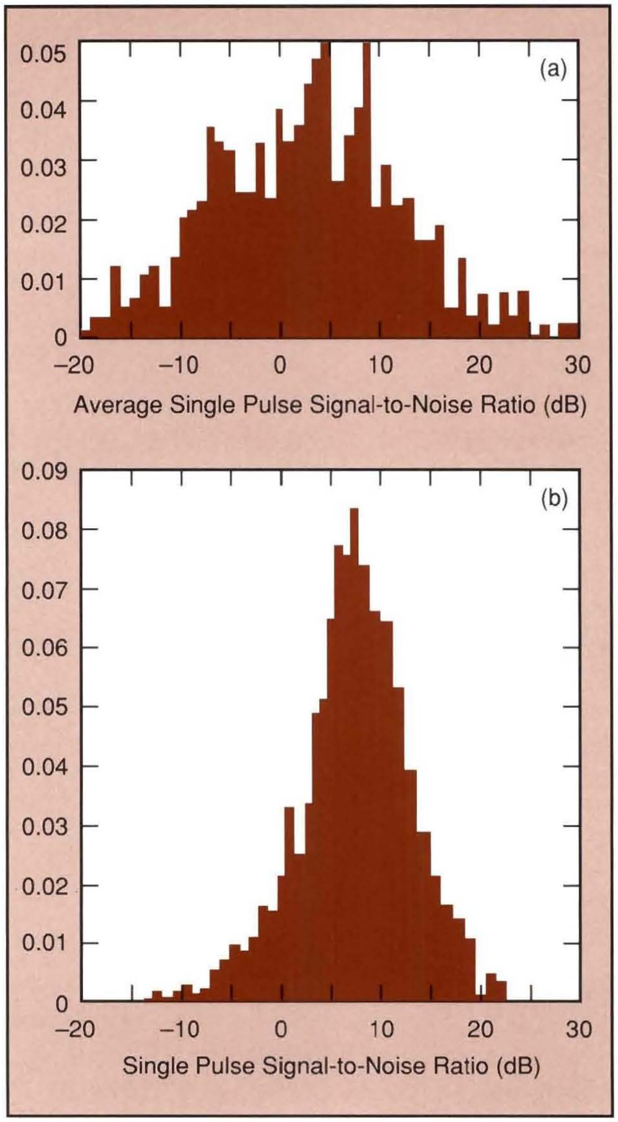

Fig. 1—Histograms that characterize the typical high-earth-orbit (HEO) satellite tracking environment. (a) The fraction of 722 targets per dB of SNR versus average single-pulse SNR. Because more than one-half of the area under the histogram curve is associated with SNRs below a minimum detection threshold of 10 dB, most targets cannot be detected by the sensor on the basis of a single radar pulse. Hence multiple-pulse detection and estimation signal processing play an important role in the design of a HEO satellite surveillance radar. (b) The fraction of the total number of pulses per dB of SNR versus the single-pulse SNR value during a single track of a single target. More than 30 dB of SNR fluctuation occurs within this track. Hence the multiple-pulse detection and estimation processing models must allow for fluctuating target cross section.

The histogram in Fig. 1(b) characterizes an ensemble of single-pulse SNRs obtained during a one-minute track of a single target. The fraction of the total number of pulses per dB of SNR is plotted versus the single-pulse SNR value. Because all other relevant variables are approximately constant over this track time, the histogram shows that target cross-section fluctuation is more than 30 dB. The histogram cannot convey how much amplitude variation occurs within a smaller set of multiple pulses; in the worst case the amplitude variation could be the full dynamic range of the target track.

In some radar applications, wide signal bandwidths and high data rates make the design problem difficult. In the narrowband ground-based HEO satellite tracking radar the combination of low SNRs and fluctuating cross section determines the signal processing design. The signal processing models must include provisions for both multiple-pulse and fluctuating cross-section features. These requirements influence the models for signal amplitude and delay that are related to the physically intuitive notions of target cross section and range, hence their importance to the satellite tracking problem. The radar receiver signal amplitude is proportional to the target cross section, so a complex variation of signal amplitudes is often observed in a small given set of multiple-pulse data because of target cross-section fluctuation. The signal delay (the time required for the radar signal to propagate to the target and return to the radar) is proportional to the target range, so a smooth variation in delay is often observed in a small given set of multiple-pulse data because of changing target range.

The Filter-Bank Receiver

Radar signal processing is often considered in the context of the analog hardware architec-

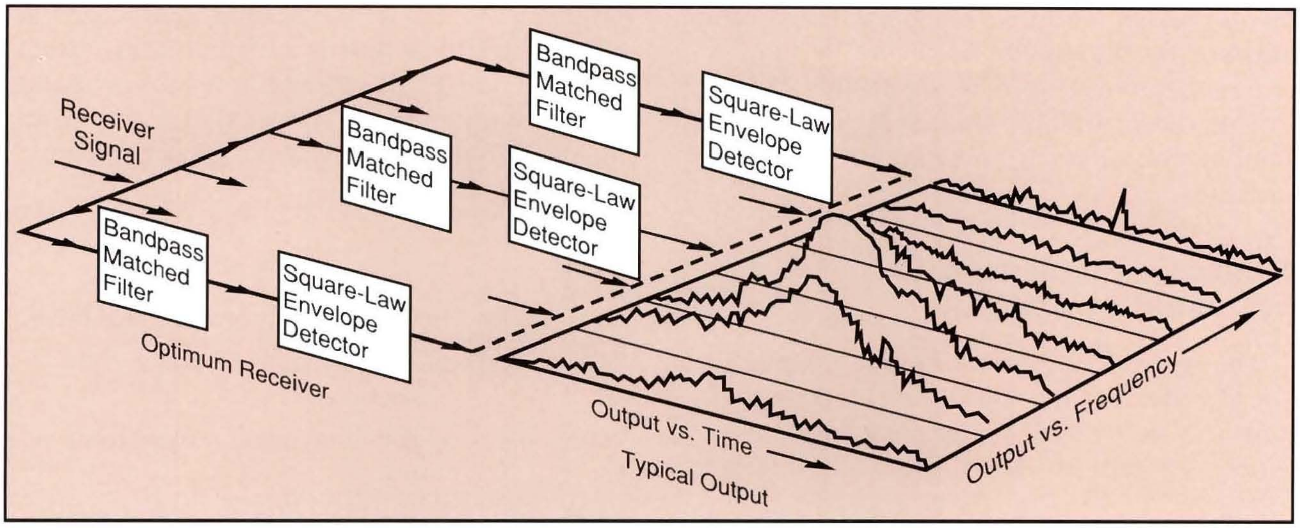

Fig. 2—The filter-bank receiver for a slowly fluctuating point target. This receiver contains the linear (bandpass matched filter) and nonlinear (square-law envelope detector) components common to most optimal designs. Similar receivers were originally built with dedicated analog hardware. The signal processing model associated with this receiver needs to be generalized for the HEO satellite tracking problem (drawing modified from Van Trees [1], p. 278).

tures used before the arrival of digital signal processors and general-purpose digital computers. An interesting example of this viewpoint is the discussion of slowly fluctuating point-target models found in Chapter 10 of Ref. 1. Figure 2 shows a filter-bank radar receiver designed to be used against a slowly fluctuating point target.

The slowly fluctuating point-target model is not as general as the model considered in this article for HEO satellite tracking, but the radar receiver designed to track such targets illustrates several important points about radar signal processing system design over the last few decades. In the frequency domain, a change in range between the target and the radar introduces a frequency, or Doppler, shift in the radar receiver signal, relative to the frequencies transmitted by the radar. In the time domain this effect shows up as a time-varying signal delay. In the filter-bank receiver, the N filter elements in the bank correspond to N possible frequency shifts. If the frequency shift cannot be modeled as a constant because of higher-order range changes between the target and radar, then the filter bank would need at least N^2 filter elements to correspond to N possible frequency shift rates for each of N possible frequency shifts. Thus, as the order of the Doppler-shift model is increased, the necessary number of filter-bank elements increases exponentially. It does not matter whether the problem is cast in the time or frequency domain. Multiple-pulse models often include higher-order Doppler-shift models or delay models to represent a complicated variation in target range. This article generalizes the receiver design shown in Fig. 2 by estimating higher-order range rates with a method that has a computational advantage over a fully generalized filter-bank implementation.

A unified explanation of the successful radar receiver designs implemented in radar's recent history is difficult to find in any reference. Figure 2 shows one of the common elements found among optimal designs—the receiver that contains linear subcomponents performs a nonlinear transformation on the signal. Because the bandpass matched filter is a linear process performed on the receiver signal, and the square-law envelope detector represents a purely nonlinear process, the overall transformation is nonlinear. The method of estimating signal amplitude and delay proposed in this article also decomposes into linear and nonlinear components. Receivers similar to the model in Fig. 2 were originally built with dedicated analog hardware. Many modern algorithms that are implemented as software on a general-purpose digital computer merely emulate the old analog designs. When the constraints imposed by the analog hardware designs are removed, interesting new algorithms emerge. One new algorithm is obtained by studying the solutions of nonlinear problems with separable variables.

Nonlinear Problems with Separable Variables

The signal-amplitude and delay model implied by the matched-filter-bank receiver of Fig. 2 for the slowly fluctuating point target is not general enough for the HEO satellite tracking problem. The model needs to be generalized to allow for multiple pulses and fluctuating target cross section. Figure 3 illustrates a narrowband linear channel/target model that leads to an ideal baseband receiver signal [2] of the form

$$s(\alpha,\tau)(t) \equiv \alpha(t)x\big(t-\tau(t)\big)v(t) \tag{1}$$

where

$\alpha : t \to \alpha(t)$ is the signal-amplitude function,

$x : t \to x(t)$ is the transmitter-baseband function,

$\tau : t \to \tau(t)$ is the true signal-delay function, and

$v : t \to v(t)$ is the local-oscillator function.

It is easy to verify by substitution of the definition in Eq. 1 that

$$\begin{aligned} s(a\alpha+\beta,\tau)(t) &= \big(a\alpha(t)+\beta(t)\big)x\big(t-\tau(t)\big)v(t) \\ &= a\alpha(t)x\big(t-\tau(t)\big)v(t)+\beta(t)x\big(t-\tau(t)\big)v(t) \\ &= as(\alpha,\tau)(t)+s(\beta,\tau)(t). \end{aligned} \tag{2}$$

Equation 2 shows that the receiver signal defined by Eq. 1 is linear in the amplitude variable α but not necessarily linear in the delay variable τ. The solution for α and τ, given some observation of $s(\alpha,\tau)$, is a nonlinear problem with separable variables. Thinking of α and τ as functions instead of scalars can accommodate some complex signal-amplitude and time-delay models. A particular case corresponding to the HEO satellite tracking problem is derived in the appendix.

The structured nonlinear problem with separable variables was studied in the mid-1970s by G.H. Golub and V. Pereyra [3], and F. Krogh [4], and later by L. Kaufman and Pereyra [5], and A. Ruhe and P.A. Wedin [6]. The original problem of solving for α and τ can be reformulated into two problems—one that depends linearly only on α and another that depends (possibly nonlinearly) only on τ. This reformulation, which is important because linear problems are well understood and efficient algorithms exist for their solution, leads to an effective technique for estimating the signal amplitude and delay.

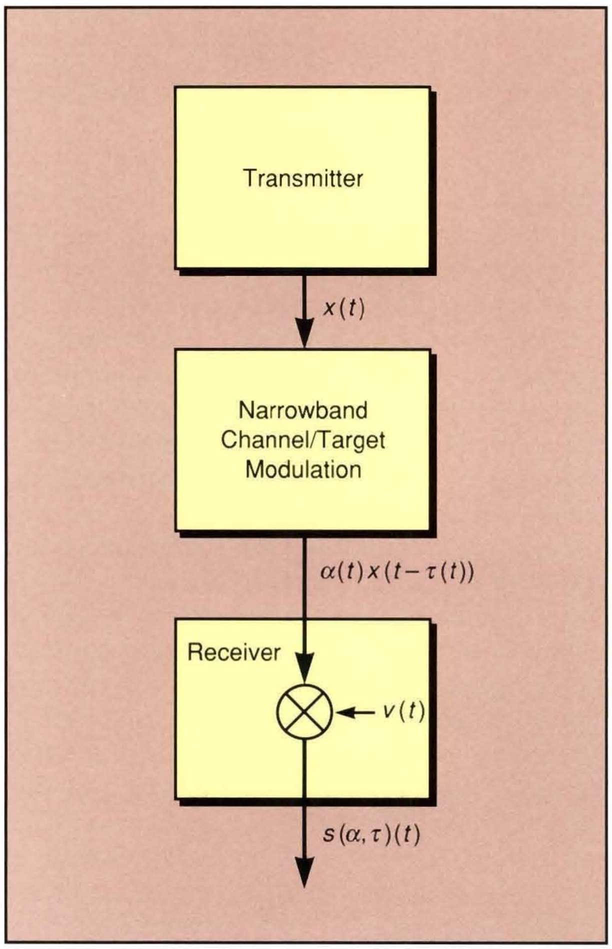

Fig. 3—A narrowband linear channel/target model. Signal processing models suitable for narrowband radars lead to a nonlinear signal processing problem with separable variables. This special structure permits an efficient algorithm for estimating radar signal amplitude and delay, which are then converted to estimates of target cross section and range.

Best Approximation of Signal Amplitude and Delay

This section discusses the estimation of receiver signal amplitude and delay from the best-approximation, or minimum-norm, viewpoint in a Euclidean space [2]. From this viewpoint many results can be derived as consequences of well-known principles (such as the projection theorem and the Pythagorean relationship). Thus detailed proofs, which are seldom intuitive, can be avoided.

The best approximation of receiver signal amplitude and delay is more general than the particular estimation problem encountered in HEO satellite tracking. New transmitter-baseband waveforms and signal-amplitude functions in Eq. 1 provide the variations needed to address other radar signal processing problems. For example, some sensors use phase-coded transmitter functions instead of the frequency-coded function assumed in the appendix; this modification is easily made.

Best approximation is one viewpoint that naturally leads to an algorithm architecture that can be implemented efficiently with modern computational methods. The algorithm architecture exploits linearity in a manner that does not obviously emulate the analog radar receiver architectures used for so many years. Hence a new way of thinking about the receiver design problem is exposed.

The Best-Approximation Problem

This section defines what is meant by the approximating set, ideal receiver signals, and an observed receiver signal. A few definitions establish the underlying structure of the problem. Let H be a Hilbert space of complex-valued functions of the real (time) line with the usual pointwise operations of addition and scalar multiplication. The functions in H can be called signals, elements, or points of H, and all of these terms are interchangeable in this article. In the space H the norm is denoted by $\|\cdot\|$ and the inner product by $\langle\cdot,\cdot\rangle$. Let A be a subspace of H and let T be some subset of H. The subspace A contains the functions that represent possible receiver signal amplitudes, and the subset T contains functions that represent possible signal delays. The *ideal receiver signal* is denoted by a function $s: A \times T \to H$ that satisfies the following property:

PROPERTY 1 (SEPARABLE VARIABLES). *The function $s(a\alpha + \beta, \tau) = as(\alpha, \tau) + s(\beta, \tau)$ for all complex scalars a, α and β in A, and τ in T.*

Each ideal receiver signal is built from an amplitude function and a delay function. A specific ideal receiver signal model (including the amplitude and delay models) for the HEO satellite tracking problem is developed in the appendix.

The description above is all of the structure needed to study the best approximation of an element of H (some observed radar receiver signal) on the *approximating set*

$$S \equiv \left\{ s(\alpha, \tau) \middle| \, \alpha \text{ in } A, \tau \text{ in } T \right\}$$

of ideal receiver signals. The problem is to find a noiseless ideal radar receiver signal from S that best approximates, in a minimum norm sense, some noisy *observed radar receiver signal* represented by a general element of H.

The remainder of this section defines the best-approximation problem and illustrates some methods to solve it. Once the best-approximating ideal receiver signal is found, the associated signal-amplitude and delay functions are also recovered through Eq. 1, and they determine the desired estimates of signal amplitude and delay.

We introduce a particular subset of S before defining the best-approximation problem. The set S_τ is defined for each τ in T in terms of ideal receiver signals by

$$S_\tau \equiv \left\{ s(\alpha, \tau) \middle| \, \alpha \text{ in } A \right\}.$$

Property 1 implies that for each τ in T, S_τ is a subspace of H. This fact is important because best approximation on a subspace is a well-structured problem. Note also that

$$S = \bigcup_{\tau \text{ in } T} S_\tau$$

so that the union of all of these subspaces over all possible values of τ regenerates the approximating set S.

The best approximation of the observed receiver signal r in H on the set S is the ideal receiver signal $\hat{r}$ in S such that

$$\|r - \hat{r}\| \leq \|r - y\|$$

for all y in S. Let us consider the following questions: (1) when does the best approximation exist, (2) when is it unique, and (3) how is it characterized, or constructed? The third question—the construction of the best approximation—is the focus of this article; the issues of existence and uniqueness are not directly addressed here.

Construction of the Best Approximation

Four problems are defined in this section. Each problem is either a best-approximation problem or is closely related to a best-approximation problem. We assume that each problem has at least one solution.

The best approximation of signal amplitude and delay is the problem of interest defined by

PROBLEM 1 (PRINCIPAL). *Given r in H, find a best approximation of r on S; that is, find $\hat{\alpha}$ in A and $\hat{\tau}$ in T such that*

$$\|r - s(\hat{\alpha}, \hat{\tau})\| \leq \|r - s(\alpha, \tau)\|$$

for all α in A and τ in T.

The function $s(\hat{\alpha}, \hat{\tau})$ is the solution of the principal problem. By a well-known projection theorem [7] the error $r - s(\hat{\alpha}, \hat{\tau})$ is orthogonal to the subspace $S_{\hat{\tau}}$. In particular,

$$\langle r - s(\hat{\alpha}, \hat{\tau}), s(\alpha, \hat{\tau}) \rangle = 0$$

for all α in A. The function $s : (\alpha, \tau) \to s(\alpha, \tau)$ of two variables is generally nonlinear in τ. If the function were linear, then many efficient methods would be available to solve the principal problem for signal amplitude and delay. A first step in solving such a nonlinear problem is to reduce the problem to finite dimensions. This reduction prepares the problem for a numerical solution that must be done in a finite dimensional space. A reduction in the dimensionality is accomplished by introducing a finite-dimensional parameterization of the sets A and T. A finite-dimensional parameterization of the set A is a function $f : P_A \to A$ from the finite-dimensional set P_A onto the set A. Since $f : P_A \to A$ denotes a finite-dimensional parameterization of the set A, let $g : P_T \to T$ denote a finite-dimensional parameterization of the set T. Figure 4 is a diagram that illustrates the point, set, and functional relationships associated with the best-approximation problem.

The most obvious numerical solution of the general nonlinear best-approximation problem is to minimize the function $\|r - s(\alpha, \tau)\|$ on $P_A \times P_T$, or the set of all pairs consisting of one element from P_A and one element from P_T. This solution yields a set of parameter values that minimizes the norm and determines the estimate of signal amplitude and delay. Unfortunately, a direct numerical approach is only practical with a few unknown parameters. In particular, for the HEO satellite tracking problem the number of amplitude parameters (the dimension of P_A) is too large.

Instead of solving the principal problem as a general nonlinear problem, we can use the separable-variable property to reduce the dimension of the numerical problem [2–6]. The principal problem thus contains

PROBLEM 2 (LINEAR). *Given r in H and τ in T, find a best approximation of r on S_τ; that is, find $\hat{\alpha}_\tau$ in A such that*

$$\|r - s(\hat{\alpha}_\tau, \tau)\| \leq \|r - s(\alpha, \tau)\|$$

for all α in A.

The function $s(\hat{\alpha}_\tau, \tau)$ is the solution of the linear subproblem. In this case, the projection theorem guarantees that the error $r - s(\hat{\alpha}_\tau, \tau)$ is orthogonal to the subspace S_τ. In particular,

$$\langle r - s(\hat{\alpha}_\tau, \tau), s(\alpha, \tau) \rangle = 0 \qquad (3)$$

for all α in A. This problem is simpler than

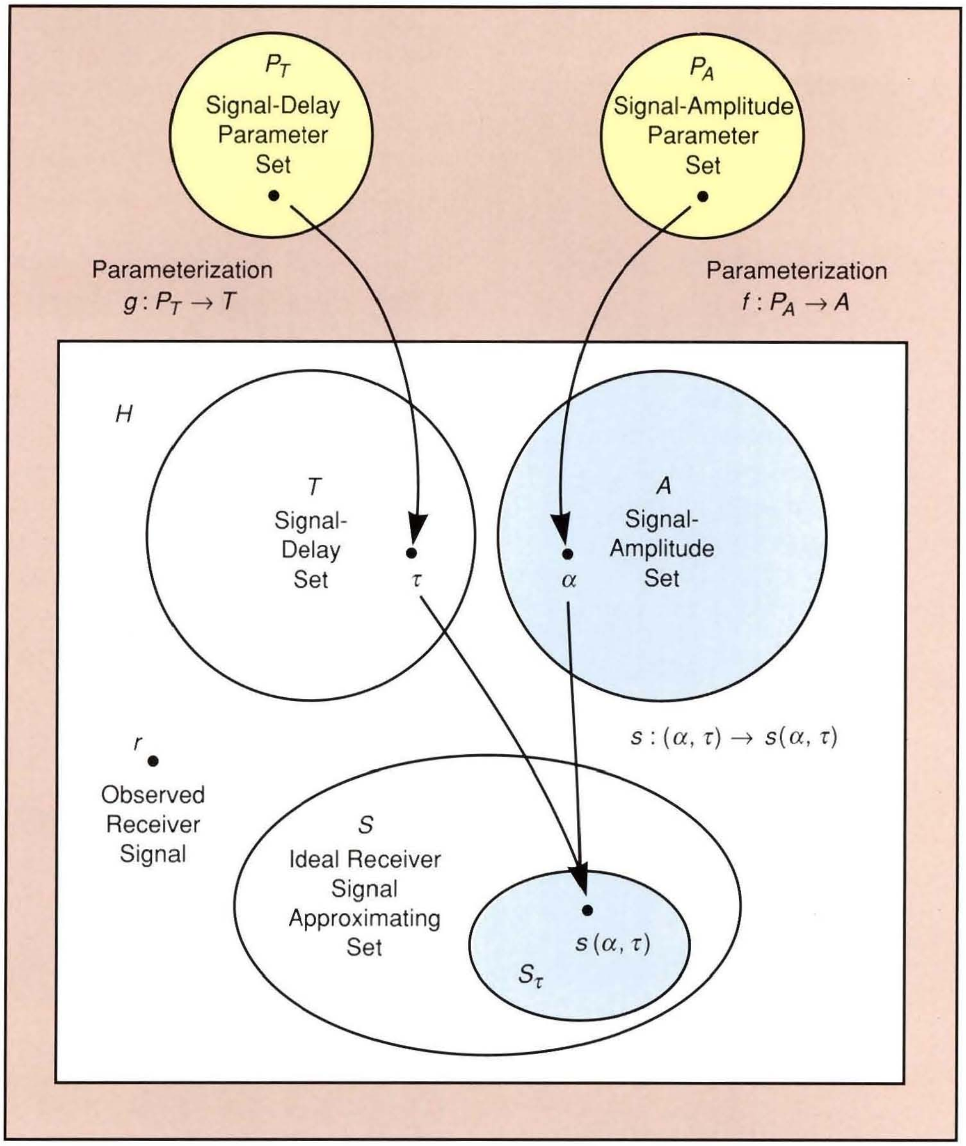

Fig. 4—The point, set, and functional relationships used in the mathematical description of the best approximation of radar signal amplitude and delay. This diagram shows subspaces as blue sets. A minimum norm (minimum error) approximation of the observed receiver signal on the approximating set S of ideal receiver signals is associated, through the modeling functions, with a best set of signal amplitude and delay parameter estimates.

Problem 1 because the approximating set is now S_τ rather than S, and S_τ is a subspace of H. Best approximation on a subspace is a standard linear problem, and the linear subproblem can be solved efficiently by a number of techniques that are not suitable for the nonlinear problem. The number of unknowns in the parameterization of A (the dimension of P_A) that can be practically accommodated in Problem 2 is also larger than the number of unknowns that can be treated in a general nonlinear problem. The appendix contains additional remarks that concern a closed form solution of the linear subproblem for the HEO satellite tracking problem.

To use the solution $s(\hat{\alpha}_\tau, \tau)$ of the linear subproblem to construct a solution of the principal problem, define

PROBLEM 3 (ALTERNATE). *Given r in H, find $\hat{\tau}$ in T such that*

$$\left\| r - s(\hat{\alpha}_{\hat{\tau}}, \hat{\tau}) \right\| \le \left\| r - s(\hat{\alpha}_\tau, \tau) \right\|$$

for all τ in T.

Thus $s(\hat{\alpha}_{\tilde{\tau}}, \tilde{\tau})$ denotes the solution of the alternate problem, which is also the solution of the linear subproblem for a particular value of τ.

A simple inequality chain

$$\left\|r - s(\hat{\alpha}_{\tilde{\tau}}, \tilde{\tau})\right\| \le \left\|r - s(\hat{\alpha}_{\tau}, \tau)\right\| \le \left\|r - s(\alpha, \tau)\right\|$$

that holds for each α in A and τ in T implies that the solution of the alternate problem also solves the principal problem. Thus we construct a solution to the principal problem by numerical optimization over T, which utilizes the closed-form solution of the linear subproblem.

The alternate problem is not the only nonlinear optimization problem over T that also solves the principal problem. Suppose $\tilde{\tau}$ in T satisfies

$$\left\|s(\hat{\alpha}_{\tilde{\tau}}, \tilde{\tau})\right\|^2 \ge \left\|s(\hat{\alpha}_{\tau}, \tau)\right\|^2 \tag{4}$$

for all τ in T. The orthogonality condition in Eq. 3 allows the Pythagorean property to be applied to both sides of

$$\|r\|^2 - \left\|s(\hat{\alpha}_{\tilde{\tau}}, \tilde{\tau})\right\|^2 \le \|r\|^2 - \left\|s(\hat{\alpha}_{\tau}, \tau)\right\|^2,$$

which follows from the inequality in Eq. 4. Then

$$\left\|r - s(\hat{\alpha}_{\tilde{\tau}}, \tilde{\tau})\right\|^2 \le \left\|r - s(\hat{\alpha}_{\tau}, \tau)\right\|^2 \tag{5}$$

for all τ in T. If we compare the inequality in Eq. 5 to the alternate problem, it follows that a solution of the alternate problem (and hence the principal problem) can also be found by solving

PROBLEM 4 (FINAL). *Given r in H, find $\tilde{\tau}$ in T such that*

$$\left\|s(\hat{\alpha}_{\tilde{\tau}}, \tilde{\tau})\right\| \ge \left\|s(\hat{\alpha}_{\tau}, \tau)\right\|$$

for all τ in T.

Thus the principal problem can be solved by each of the following methods:

1. Utilize the solution of the linear sub-problem (Problem 2) to solve the alternate best-approximation problem (Problem 3) by numerical optimization over P_T.
2. Utilize the solution of the linear subproblem (Problem 2) to solve the final problem (Problem 4) by numerical optimization over P_T.
3. Solve the principal best-approximation problem (Problem 1) by numerical optimization over $P_T \times P_A$.

Many interesting cases exist in which a closed-form solution of the linear subproblem is easily computed, as in the HEO satellite tracking problem described in the appendix. In these cases, if we use methods 1 or 2, we avoid numerical optimization over the set P_A of parameters that characterize A. Thus computation based on methods 1 or 2 may be considerably faster than computation based on method 3.

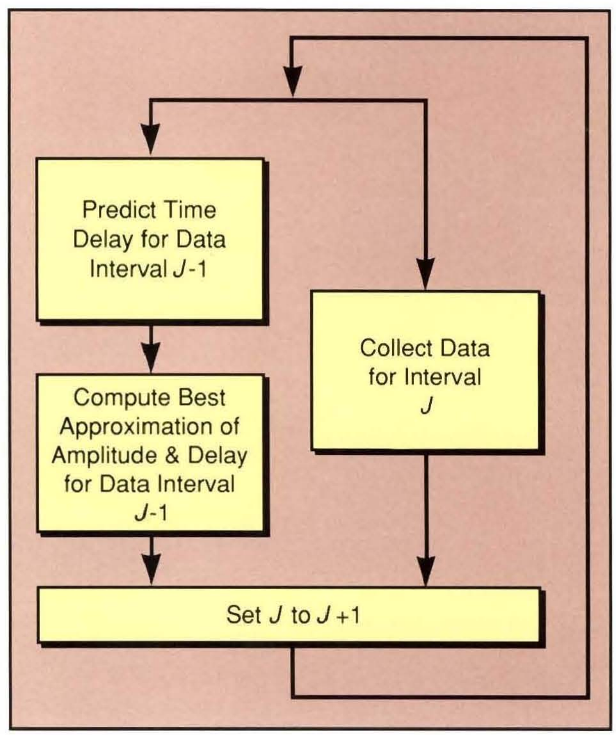

Fig. 5—Direct estimation of signal amplitude and delay. A practical algorithm is obtained by combining the best approximation of radar signal amplitude and delay with delay prediction and data collection steps. Delay prediction is required to initialize the numerical optimization process (used within the best-approximation method) on one set of data while the next data set is being collected.

Implementation and Testing

The previous section described the problem of estimating the radar receiver signal amplitude and delay from the best-approximation viewpoint. The signal amplitude is proportional to the amount of energy reflected from the target. The signal delay (the time required for the radar signal to propagate to the target and return to the radar) is proportional to the target range. The signal amplitude and delay can be converted to the physically intuitive notions of target cross section and range at any time, hence their importance to the satellite tracking problem.

The model that characterizes the signal amplitude α and delay τ is developed in detail in the appendix. The form of this model is suitable for the HEO satellite tracking problem. The signal amplitude is constant for the duration of a radar pulse but varies unrestricted from pulse to pulse. Parameters in the delay model represent the values of the delay and its rates at a reference time. If enough rate parameters are included, the delay model can describe any smooth variation of delay with time.

Figure 5 illustrates how direct estimation of delay can be accomplished by combining an implementation of one of the solutions of the best-approximation problem with delay prediction and data collection steps. The prediction step provides a crude starting point for the numerical maximization of $\|s(\hat{\alpha}_\tau, \tau)\|$ over all possible values of τ; this starting point is required to solve the final problem (Problem 4) and to estimate the best approximation of signal amplitude and delay. We implement the data collection step in parallel with computation, so that the next interval of data is available when computation on the previous data interval is completed.

In steady-state operation, the delay is first predicted from the current state of the target dynamics model. The target dynamics model is chosen to reflect what the designer knows about the behavior of the target. In the HEO satellite tracking problem an orbital dynamics model is appropriate. The predicted values of delay provide the initial value of τ for the numerical optimization routine that maximizes $\|s(\hat{\alpha}_\tau, \tau)\|$. The numerical maximization of $\|s(\hat{\alpha}_\tau, \tau)\|$ requires repeated evaluation of $\|s(\hat{\alpha}_\tau, \tau)\|$ with different values of τ. Each evaluation involves a linear transformation of the radar receiver signal data, which solves the linear subproblem (Problem 2). The appendix on p. 325 describes the specific transformation for the special case related to HEO satellite tracking. The value of τ obtained by numerical maximization is used to update the target dynamics model, and the entire process is repeated with the next set of radar receiver signal data.

Numerical Optimization

To assess the practicality of implementing the signal-amplitude-estimation and delay-estimation algorithms illustrated in Fig. 5, a simplified version of an optimization routine due to R.P. Brent [8] and the transformation required for solving the linear subproblem (Problem 2) were coded for a Floating Point Systems AP120B processor. The optimization algorithm was essentially identical to the algorithm that Golub and Pereyra used in their numerical studies of the solution of nonlinear problems with separable variables [3].

The optimization routine performs successive quadratic approximations to estimate the maximum value of $\|s(\hat{\alpha}_\tau, \tau)\|$. Figure 6 shows a typical example for a linear delay model, which is obtained by setting $M = 1$ in the more general model of Eq. A3 developed in the appendix. Figure 6 shows the curves determined by equal values of $\|s(\hat{\alpha}_\tau, \tau)\|$ in τ space. The algorithm progresses from the predicted delay at point A to the vicinity of the best approximation of delay at point B in 12 optimization steps. Each step requires one solution of the linear subproblem. In the example shown the predicted delay was 0.42 μsec (63 m) and 0.40 μsec/sec (60 m/sec) from the best approximation.

Sources of Estimation Error

Information about the errors between the value of τ obtained by numerical optimization of $\|s(\hat{\alpha}_\tau, \tau)\|$ and the true value of τ (the true value

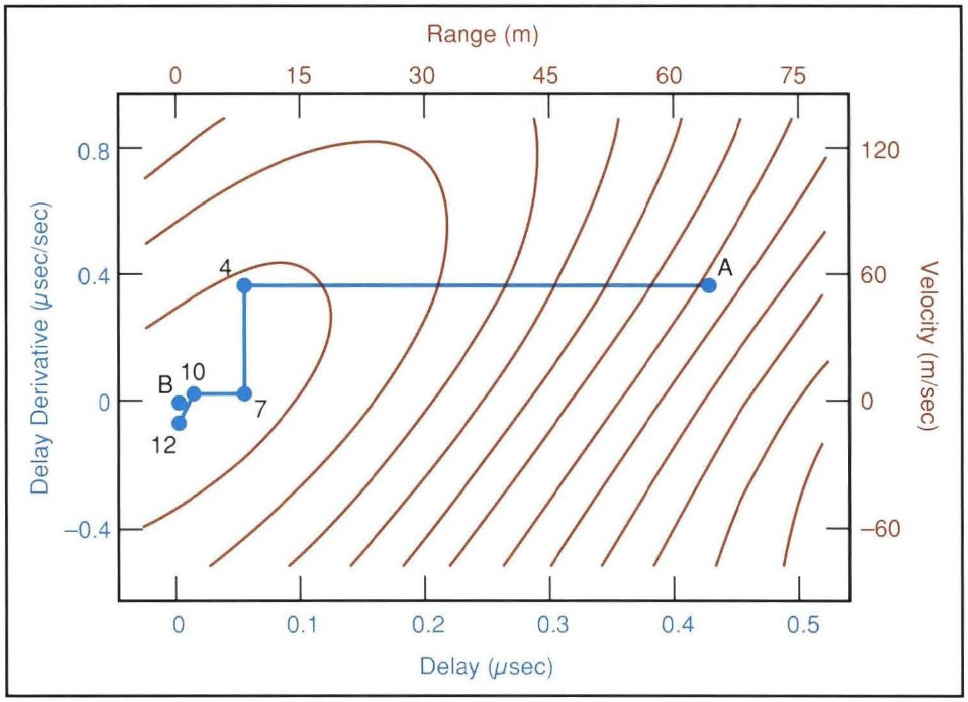

Fig. 6—A typical example of the numerical optimization process. Numerical optimization reduces the large error between the predicted delay for the data set (labeled A) and the best approximation of delay (labeled B). The best approximation of delay is the value that maximizes $\|s(\hat{\alpha}_{\tau}, \tau)\|$. The curves are defined by equal values of $\|s(\hat{\alpha}_{\tau}, \tau)\|$ in τ space. The numbered points show the path of the optimization process over the 12 steps required to progress from the predicted delay to the vicinity of the best approximation of delay. The best approximation of signal amplitude is determined by the solution of the linear subproblem associated with the best approximation of delay.

of delay) is needed to update the target dynamics model. The value of the best approximation $\hat{\tau}$ is not identical to the true value of delay because of noise and other errors that corrupt the observed radar receiver signal. The fact that any real implementation can compute only a numerical approximation of $\hat{\tau}$ is often overlooked. Thus an error occurs in the estimate because noise corrupts the observed receiver signal (the difference between the best approximation $\hat{\tau}$ and the true value of delay), and an additional error occurs because of the numerical error of the algorithm (the difference between $\hat{\tau}$ and the numerical approximation of $\hat{\tau}$).

Many computational problems produce numerical errors due only to roundoff error in the computer; this fact is not always true for numerical optimization problems. Theoretically, an infinite number of optimization steps is required to compute the maximum value of the function being optimized, unless the function is quadratic or satisfies some rare special cases. In practice, optimization must stop after reasonable execution time. Terminating the optimization process too early may cause the final numerical error to be larger than the roundoff error.

The best-approximation approach to the design of an estimator of signal amplitude and delay does not require any assumptions concerning how the observed radar receiver signals differ from the ideal receiver signals. Usually the ideal receiver signals represent all possible noiseless radar receiver signals; the actual observed signals differ from these ideal signals only by an additional additive noise component. The assumptions about the noise determine the signal-amplitude-estimate and delay-estimate errors, provided that the algorithm includes accurate models of the radar and the target behavior (see the box titled "Delay Estimate Errors due to Observation Noise").

How well does the method and its implementation numerically estimate the best approximation of signal amplitude and delay? In the HEO satellite tracking problem, the estimation of delay is more important than the estimation of amplitude. The delay estimates are needed to compute the position measurements used to determine the satellite's orbit.

Numerical Tests

A series of tests was performed to determine if the algorithms could consistently estimate the best approximation of delay associated with tracking a HEO satellite. This section presents the results from one of those tests. The test was designed to measure the typical numerical error encountered in estimating the best approximation $\hat{\tau}$ of delay. The numerical error includes the combined errors due to termination of the optimization procedure and the roundoff error. The test used a quadratic delay model obtained by setting $M = 2$ in the more general model developed in the appendix (Eq. A3). Table 1 summarizes the relevant radar and signal parameters. The algorithm resolved 10,000 random errors between the predicted delay that was used to initialize the numerical optimization routine and the best approximation of delay.

Table 1. Radar and Signal Parameters

Interpulse Period: 0.036 sec
Pulse Width: 1.04 msec
Bandwidth: 500 kHz
Number of Trials: 10,000
Processing Time: 60% of Real Time
Per Pulse S/N: 0 dB
Pulses Per Processing Interval: 200

Figure 7 shows histograms of the range and range-rate errors determined from the delay errors. Two histograms are shown for each target range parameter, along with sample standard deviations. In the three-dimensional problem the three range-error parameters are

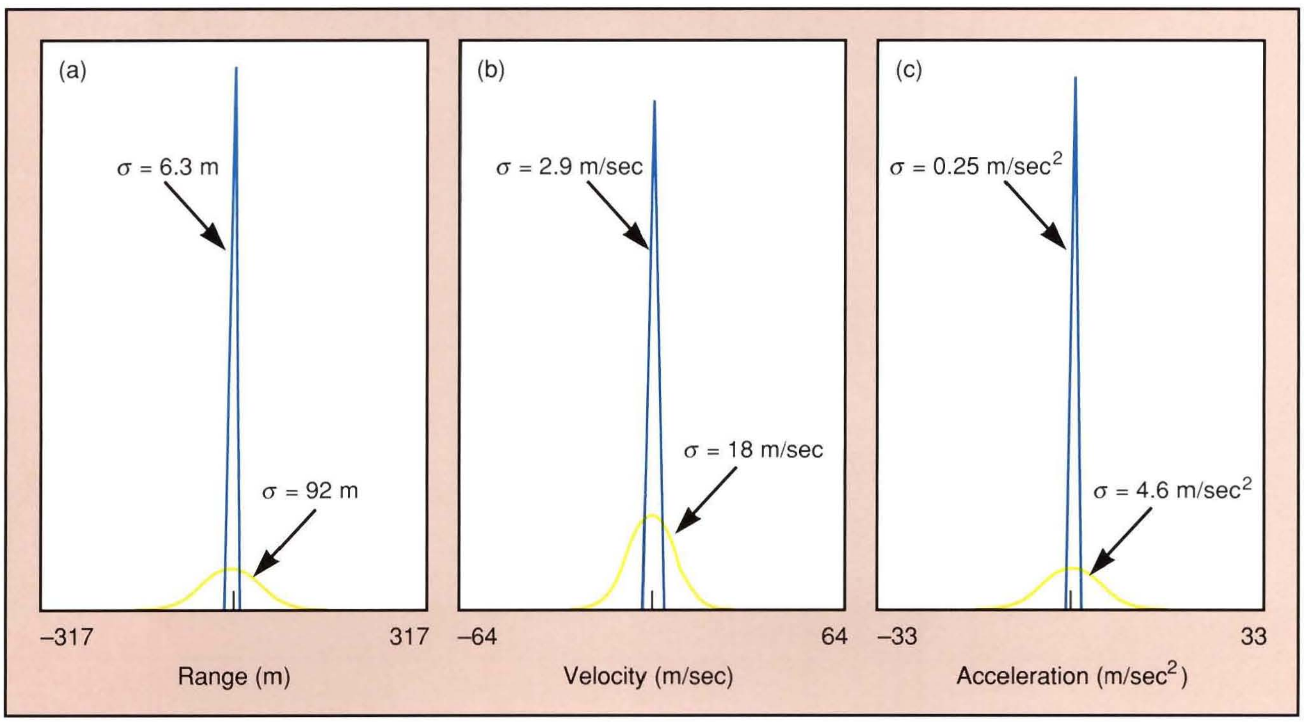

Fig. 7—Determination of numerical optimization errors in target range, velocity, and acceleration. Each pair of histograms shows the error before and after the numerical optimization process. An order-of-magnitude reduction in the standard deviation of the errors is obtained. The final errors are small enough so that the total range error is dominated by noise errors, and not the numerical error of the optimization process.

Delay-Estimate Errors due to Observation Noise

A well-designed implementation of an estimator is limited only by the effects of noise in the observed radar receiver signal, and not by numerical errors due to computer roundoff and termination of numerical optimization. In the signal-amplitude-estimation and delay-estimation problem for HEO satellite tracking, components of noise arise from sky noise (atmospheric, cosmic background, galactic, and solar) and radar receiver noise (shot and thermal). The effect of clutter (backscatter from objects other than the desired target) is often treated as noise, but it is usually negligible in the HEO satellite tracking problem.

From a stochastic model of the observed receiver signal, the effects of noise on estimate error can be calculated by using statistical estimation theory [1]. These techniques yield explicit results for the signal-amplitude-estimate and delay-estimate errors due to noise. Because of the characteristics of the dominant noise encountered in the HEO satellite tracking problem, a reasonable model for the observed receiver signal consists of a complex Gaussian stochastic process added to the ideal receiver signal. The stochastic process is assumed to be stationary and white. A realization of the observed stochastic receiver signal is then defined as

$$r(t) \equiv s(\alpha, \tau)(t) + w(t)$$

where $s(\alpha, \tau): t \to s(\alpha, \tau)(t)$ is the ideal receiver signal and $w: t \to w(t)$ is a realization of the noise. The ideal receiver signal is determined by the amplitude function $\alpha: t \to \alpha(t)$ and the delay function $\tau: t \to \tau(t)$ from Eq. 1.

With these assumptions, the estimates $\hat{\alpha}$ and $\hat{\tau}$ that result from the best approximation of signal amplitude and delay are also the maximum likelihood (ML) estimates. That is, the pair $(\hat{\alpha}, \hat{\tau})$ maximizes (over the set $A \times T$) the joint probability density of the observed receiver signal evaluated at its realization r.

The error in the ML estimates due to noise is described by the Cramer-Rao (C-R) bound. Any unbiased estimator has error covariances no smaller than the C-R bound, but the C-R bound has special significance for ML estimators, unbiased or not. In a multiple-pulse implementation, a total SNR is defined as the sum of the single-pulse SNRs. Under reasonable conditions, as the total SNR increases, the distribution of the errors of the ML estimator tends toward a Gaussian distribution with zero mean and a covariance equal to the C-R bound. Thus the C-R bound approximates the mean-squared errors of such an estimator, and the approximation becomes exact as the total SNR becomes large.

The appendix outlines the parameterization of the signal-delay function that is used in the HEO satellite tracking problem. The three graphs of Fig. A show the error in the estimates of these parameters derived from the C-R bounds. This particular case corresponds to the conditions described in the section titled "Numerical Tests" on p. 321. The errors are presented as physically intuitive target-range, velocity, and acceleration estimate errors. The ordinates of each graph are proportional to the value of the square root of the appropriate diagonal element of the C-R bound matrix. The graphs represent the approximate rms errors in the parameter estimates due to noise; they show the well-known increase in range estimate accuracy as a function of increasing radar bandwidth. The graphs also show the relative insensitivity of velocity and acceleration estimate accuracy to radar bandwidth.

The total SNR for this example is a moderate 23 dB. Sample rms errors (as determined by comparing the delay estimates to precisely fitted orbits from laser radar sensor data) tend to be from one to three times larger than these C-R bound values, depending on the total SNR. These orbits are accurate enough to be considered exact for purposes of determining the signal-delay-estimate errors achieved on a portion of a single satellite track by a narrowband sensor. Based on such comparisons, actual rms errors due to noise in this example are typically twice the values predicted by the C-R bound.

expressed as the range, velocity, and acceleration errors, relative to the best approximation of delay associated with the track of the HEO satellite at a reference time. For each pair of histograms, the histogram with the larger standard deviation characterizes the error in the parameter before the algorithm is applied to the radar data. Thus this histogram is associated with the difference between the initial value of delay used by the optimization algo-

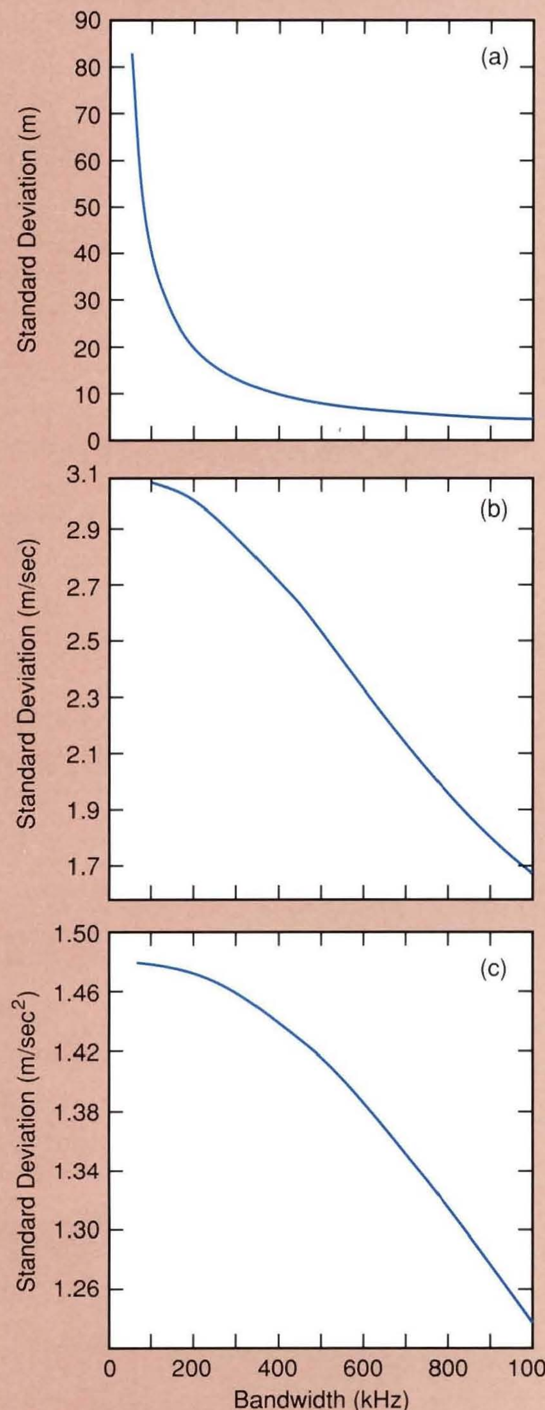

Fig. A—Cramer–Rao bounds for the standard deviation of errors in (a) target range, (b) target velocity, and (c) target acceleration, as a function of the radar transmitter bandwidth for the radar and signal parameters summarized in Table 1. The sample standard deviations encountered in a real system are larger because the C-R bounds only approximate the error due to noise in the observed radar receiver signal.

rithm and the best approximation $\hat{\tau}$ of delay. The second histogram of each pair, with the smaller standard deviation, characterizes the error in the parameter after application of the algorithm to the radar data. Thus the second histogram is associated with the difference between the best approximation $\hat{\tau}$ of delay and the numerical approximation of $\hat{\tau}$.

Optimization reduces the standard deviation of the error in each variable by an order of magnitude. The final errors after optimization are small enough so that the error between the best approximation $\hat{\tau}$ and the true value of delay (a difference primarily due to noise in the observation) dominates the total error. However, the standard deviation of final error in velocity turns out to be larger (relative to the noise errors) than the standard deviations of error in range and acceleration. Appropriate problem scaling balances each component of numerical error relative to the noise errors. Scaling is a mathematical change of delay variables that alters the multidimensional shape of $\|s(\hat{\alpha}_{\tau}, \tau)\|$ near its maximum (Ref. 9, section 8.7). Appropriate scaling transformations were added to the numerical optimization process before the final software for an operational tracking system was completed.

Conclusions

Low single-pulse SNRs and fluctuating target cross sections are encountered when ground-based radars track high-earth-orbit satellites. This condition requires that the usual filter-bank radar receiver architecture be generalized. A generalization can be obtained by estimating signal amplitude and delay in a narrowband radar from the viewpoint of best approximation in a Euclidean space. The best-approximation approach is general enough to address the problem of direct estimates of the cross section and range of the HEO satellite. In addition, this approach can be used in other specialized radar signal processing design problems by appropriately choosing waveforms and target models.

The best-approximation approach leads to a way of thinking about radar receiver design that does more than emulate the analog filter-bank architectures used for many years. With appropriate restrictions, the best-approximation approach generates estimates of target cross section and range identical to those obtained from a filter-bank receiver. Thus the algorithm can be interpreted as a generalization of the matched filter bank. For example, if a constant unknown signal amplitude is assumed, then a matched-filter-bank approach for the HEO satellite tracking problem can be constructed with generalized filter elements based on higher-order delay models. The banks can become large because the number of elements in the bank grows exponentially with the order of the delay model.

By using the algorithm proposed in this article, only the outputs of a subset of the generalized filter elements of the bank need to be constructed, and only portions of these filter outputs need to be sampled in time. This partial construction and sampling is essentially what happens whenever the linear subproblem is solved. The results of each sampling are combined with previous samplings by the numerical optimization process to predict which filter elements should be constructed and sampled next. The optimization process refines an estimate of the filter parameters that are best matched to the radar receiver signal, and the time when their outputs will maximize. This procedure is equivalent to finding the best estimate of signal amplitude and delay. The algorithm does not terminate, but the parameters of the best-matched elements, and the time of their maximum output, can be precisely predicted in a reasonable time.

The relationship between the best-approximation approach and a filter-bank receiver is less clear when a general signal-amplitude model is assumed. Thinking of the general algorithm from the best-approximation viewpoint maintains maximum flexibility when applying the algorithm to new signal processing problems. At a more philosophical level, the project experience is a reminder of the value of taking time to understand a problem from a different viewpoint. We developed an algorithm that is considerably different from a solution based on obvious extensions of a conventional filter-bank receiver.

The Millstone Hill L-band satellite surveillance radar in Westford, Mass., has used an implementation of the algorithm for several years. The authors believe the Millstone radar is the only operational radar signal processing system that uses these methods.

References

1. H.L. Van Trees, *Detection, Estimation, and Modulation Theory*, Part III (John Wiley, New York, 1971).
2. R. Raup, "Best Approximation of Signal Amplitude and Delay in a Narrowband Radar," *Technical Report 723*, Lincoln Laboratory (22 Jan. 1986), DTIC #AD-A165700.
3. G.H. Golub and V. Pereyra, "The Differentiation of Pseudo-Inverses and Nonlinear Least Squares Problems Whose Variables Separate," *SIAM J. Numer. Anal.* **10,** 413 (1973).
4. F. T. Krogh, "Efficient Implementation of a Variable Projection Algorithm for Nonlinear Least Squares Problems," *ACM Comm.* **17,** 167 (1974).
5. L. Kaufman and V. Pereyra, "A Method for Separable Nonlinear Least Squares Problems with Separable Nonlinear Equality Constraints," *SIAM J. Numer. Anal.* **15,** 12 (1978).
6. A. Ruhe and P.A. Wedin, "Algorithms for Separable Nonlinear Least Squares Problems," *SIAM Rev.* **22,** 318 (1980).
7. D. Luenberger, *Optimization by Vector Space Methods* (John Wiley, New York, 1969).
8. R.P. Brent, *Algorithms for Finding Zeros and Extrema of Functions without Calculating Derivatives* (Prentice Hall, Englewood Cliffs, NJ, 1973).
9. P. Gill, W. Murray, and M. Wright, *Practical Optimization* (Academic Press, New York, 1981).

Appendix

Reference 2 discusses how a narrowband linear channel/target model leads to an ideal baseband receiver signal $s(\alpha, \tau): t \to s(\alpha, \tau)(t)$ of the form

$$s(\alpha, \tau)(t) \equiv \alpha(t)x(t - \tau(t))v(t) \qquad \text{(A1)}$$

where

$\alpha: t \to \alpha(t)$	is the signal-amplitude function,
$x: t \to x(t)$	is the transmitter-baseband function,
$\tau: t \to \tau(t)$	is the true signal-delay function, and
$v: t \to v(t)$	is the local-oscillator function.

The local-oscillator function in a heterodyne receiver is usually chosen to control the bandwidth of $s(\alpha, \tau): t \to s(\alpha, \tau)(t)$ so that the function can be sampled at lower rates. Otherwise, the exact choice of $v : t \to v(t)$ does not affect the estimate of signal amplitude and delay. This appendix defines particular forms of the functions in the model given by Eq. A1 and outlines a solution of the linear subproblem. These particular forms were used in an implementation of the algorithm to address the HEO satellite tracking problem. In this appendix the inner product $\langle y, z\rangle$ is defined by

$$\langle y, z\rangle \equiv \int_{-\infty}^{\infty} y(t)\, z^*(t)\, dt$$

where z^* is the complex conjugate of z. Then the norm of z is given by

$$\|z\|^2 \equiv \langle z, z\rangle .$$

First we define the transmitter-baseband function. The use of complex-valued functions in signal processing often facilitates the modeling. The possibilities for the transmitter-baseband function are limited in practice by the available radar hardware. This model is for a chirp function, which is a signal with a frequency that varies linearly in time. Define $z_n : t \to z_n(t)$ as

$$z_n(t) \equiv e^{2\pi j\left(f_n t - b_n t^2/2\delta\right)}$$

to characterize the transmitter-baseband function over a single pulse. Then the complete baseband transmitter function is

$$x(t) \equiv \sum_{n=-N}^{N} p_n(t)\, z_n(t - n\Delta)$$
$$= \sum_{n=-N}^{N} p_n(t)\, e^{2\pi j\left(f_n (t-n\Delta) - b_n (t-n\Delta)^2/2\delta\right)}$$

where

$$p_n(t) \equiv \begin{cases} 1 & \text{if } -\delta/2 + n\Delta \le t \le \delta/2 + n\Delta \\ 0 & \text{otherwise.} \end{cases}$$

Here Δ and δ are real constants that define the interpulse period and the pulse width of the radar transmitter function, respectively. The function $p_n : t \to p_n(t)$ is an indicator function that defines the time support. The instantaneous bandwidth, or change in instantaneous frequency within each pulse, is $2\pi b_n$ and the center frequency of the pulse is $2\pi f_n$.

Next we define the signal-amplitude and delay functions. In the subsection titled "Construction of the Best Approximation" on p. 316, the concept of finite-dimensional parameterization of the amplitude and delay models was introduced to facilitate the numerical solution of the best-approximation problem. These parameterizations are now made explicit.

Define the set A of approximating receiver amplitude functions by assuming that the value of $\alpha: t \to \alpha(t)$ is approximately constant over a time interval corresponding to a transmitter pulse width, but can vary in an unrestricted way from pulse to pulse. This assumption is reasonable as long as the target aspect angle changes only slightly over the time corresponding to a pulse width. Thus A can be defined by

$$A \equiv \left\{ \alpha: t \to \alpha(t) \,\middle|\, \alpha(t) = \sum_{n=-N}^{N} A_n p_n(t - \tau(t)),\ A_n \text{ in } \mathcal{C} \right\} \qquad \text{(A2)}$$

where $\mathcal{C}$ is the field of complex numbers. Notice that the definition in Eq. A2 implies the existence of a parameterization from $\mathcal{C}^{2N+1}$ onto A. The values of the A_n are estimated, and the definition can be used to recover the amplitude function estimates.

Define the set T of approximating receiver delay functions by assuming that the forces acting on the target are approximately constant over the time intervals of interest. Thus the propagation delay due to the changing range between the radar and the target is a smooth function of time. The set T can then be defined by

$$T \equiv \left\{ \tau: t \to \tau(t) \,\middle|\, \tau(t) = \sum_{m=0}^{M} \tau_m t^m,\ \tau_m \text{ in } \mathcal{R} \right\} \qquad \text{(A3)}$$

where $\mathcal{R}$ is the field of real numbers. As with the definition in Eq. A2, the definition in Eq. A3 implies the existence of a parameterization from $\mathcal{R}^{M+1}$

onto T. The values of the τ_m are estimated, and the definition can be used to recover a delay function estimate.

The local-oscillator function $v : t \to v(t)$ is the last of the four functions in the model given by Eq. A1 to be specified. Like the transmitter-baseband function, the available radar hardware limits the possible forms. A common choice for the local-oscillator function is another chirp function

$$v(t) \equiv \sum_{n=-N}^{N} p_n(t - c_n) e^{2\pi j\left(\bar{f}_n(t-n\Delta-c_n)-\bar{b}_n(t-n\Delta-c_n)^2/2\delta\right)}$$

For the bandwidth of $s(\alpha, \tau): t \to s(\alpha, \tau)(t)$ to be small, it is sufficient that

$$\bar{f}_n \approx -f_n ,$$
$$\bar{b}_n \approx -b_n , \text{ and}$$
$$c_n \approx \tau(n\Delta) .$$

The values of $\bar{f}_n$, $\bar{b}_n$ and c_n must be calculated with the best available estimate of the true value of signal delay, so that $v(t) \approx x^*(t - \tau(t))$ and the above approximations are satisfied.

The complete receiver signal model is therefore

$$s(\alpha, \tau)(t) = \sum_{n=-N}^{N} A_n u_n(\tau)(t) \qquad \text{(A4)}$$

where

$$u_n(\tau)(t) = p_n(t - \tau(t)) p_n(t - c_n) e^{2\pi j \phi_n(t)} ,$$

and

$$\begin{aligned}\phi_n(t) = \; & -\bar{f}_n c_n - \bar{b}_n c_n^2/2\delta - f_n \tau(t) - b_n \tau^2(t)/2\delta \\ & + \left(\bar{f}_n + f_n + \bar{b}_n c_n/\delta + b_n \tau(t)/\delta\right)(t - n\Delta) \\ & + \left(-\bar{b}_n/2\delta - b_n/2\delta\right)(t - n\Delta)^2 .\end{aligned}$$

The radar system is designed so that receiver signal pulses do not overlap. This is accomplished by making $\delta << \Delta$. Then

$$p_n(t - \tau(t))\, p_m(t - \tau(t)) = 0 \quad \text{whenever } n \neq m.$$

It follows that

$$\langle u_n(\tau), u_m(\tau)\rangle = 0 \quad \text{whenever } n \neq m,$$

so that the set $\{u_n(\tau)\}_n$ consists of orthogonal functions.

One of the elements of the set S_τ is the solution of the linear subproblem. From Eq. A4, this solution can be written in terms of the unknown amplitude parameters $\hat{A}_n$ as

$$s(\hat{\alpha}_\tau, \tau)(t) = \sum_{n=-N}^{N} \hat{A}_n u_n(\tau)(t). \qquad \text{(A5)}$$

Because the $u_n(\tau)$ are orthogonal, the construction of the particular set $\{\hat{A}_n\}$ that solves this best-approximation problem is well known.

Clearly, for each m, $u_m(\tau)$ is the element of S_τ obtained by setting $A_n = 1$ in Eq. A4 for $n = m$ and $A_n = 0$ otherwise. Also, the orthogonality condition of Eq. 3 satisfied by the solution of the linear subproblem holds for each element of S_τ. In particular,

$$\langle r - s(\hat{\alpha}_\tau, \tau), u_n(\tau)\rangle = 0 \quad \text{for } n = -N \text{ to } N \qquad \text{(A6)}$$

where r is the observed receiver signal. Substituting Eq. A5 into Eq. A6 yields

$$\begin{aligned}\langle r, u_n(\tau)\rangle &= \sum_{m=-N}^{N} \hat{A}_m \langle u_m(\tau), u_n(\tau)\rangle . \\ &= \hat{A}_n \|u_n(\tau)\|^2 .\end{aligned}$$

Hence

$$\hat{A}_n = \frac{\langle r, u_n(\tau)\rangle}{\|u_n(\tau)\|^2} \quad \text{for } n = -N \text{ to } N. \qquad \text{(A7)}$$

Substituting Eq. A7 into Eq. A5 gives the solution of the best-approximation linear subproblem as

$$s(\hat{\alpha}_\tau, \tau)(t) = \sum_{n=-N}^{N} \frac{\langle r, u_n(\tau)\rangle}{\|u_n(\tau)\|^2} u_n(\tau)(t).$$

RICHARD C. RAUP is a senior staff member in the Surveillance Techniques Group. He received a B.S.E., an M.S.E., and a Ph.D. degree in engineering from the Johns Hopkins University. His research is in estimation theory and signal processing. Richard has been at Lincoln Laboratory since 1978, and he is a member of the IEEE, SIAM, and Sigma Xi.

ROBERT A. FORD is a staff member in the Field Systems Group. His current work focuses on real-time signal processing for wideband radar systems. Bob received B.S.E.E. and M.S.E.E. degrees from the Columbia University School of Engineering. Before joining Lincoln Laboratory, he worked for Ford Aerospace in Newport Beach, Calif. Bob is a member of Eta Kappa Nu.

GARY R. KRUMPHOLZ is a staff member in the Surveillance Techniques Group. He received a B.S. degree and a Ph.D. degree in electrical engineering from the Worcester Polytechnic Institute, and an M.A. in mathematics from the University of Wisconsin, Madison. His research speciality is detection and estimation algorithms applied to radar tracking of satellites. Gary has been at Lincoln Laboratory since 1983, and he is a member of Sigma Xi, SIAM, and the IEEE.

MARK G. CZERWINSKI is an assistant leader in the Surveillance Techniques Group. His research specialities are in space surveillance, command and control, and signal processing. He received an S.B. in electrical engineering from M.I.T. and a Ph.D. in electrical engineering from Carnegie-Mellon University.

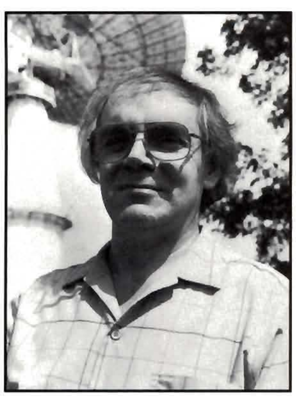

THOMAS E. CLARK is a staff member in the Surveillance Techniques Group. He received a B.A. in physics from Northeastern University, and his research speciality is satellite surveillance. Tom developed the HEO satellite tracking software used at the Millstone Hill radar in Westford, Mass. He has been at Lincoln Laboratory since 1969.