

OPTIMUM UTILIZATION OF POSITIONING DATA IN SDS III

by Gary C. GUENTHER (*) and Robert W.L. THOMAS (**)

Paper presented at HYDRO USA '86, the Second Biennial National Ocean Service International Hydrographic Conference, Norfolk, Virginia, March 1986 and reprinted with the kind permission of the National Ocean Service/NOAA and the Hydrographic Society of America.

ABSTRACT

A new, computerized hydrographic data acquisition and processing system, Shipboard Data System III (SDS III), is being designed and built for use by the National Ocean Service. An integrated positioning and navigation system is a critical element of this development. Design features include the ability to benefit from time-deskewed multiple lines of position from mixed sensor types (both electronic and manual), difficult geometries, and the use of auxiliary speed and heading data in the application of advanced filtering and smoothing techniques for reduction of random measurement noise and recognition of bias errors. Results are highly accurate, stable, and robust. Measurement noise can be reduced by as much as a factor of three without adding significant biases, even on turns, while retaining actual random vessel motions. Operations can continue during complete losses of positioning data for limited but significant periods of time, including during maneuvers.

SYSTEM DESCRIPTION

The National Ocean Service (NOS) is the sole agency responsible for the charting of the coastal waters of the United States and its territories plus the Great Lakes. This mission includes the measurement of tides and other oceanographic

(*) NOAA/National Ocean Service, 6001 Executive Blvd., Rockville, Maryland 20852, USA.

(**) EG and G Washington Analytic Services Center, Inc., 5000 Philadelphia Way, Suite J, Lanham, Maryland 20706, USA.

parameters as well as the locations of navigation aids, obstructions, landmarks, and the like. The bulk of this work is conducted by sonar from small launches in shallow water and larger ships in deep water. Shipboard Data System III is slated as the next-generation NOAA hydrographic survey data acquisition, processing, and display system for use on NOS ships and launches as replacement for the venerable HYDROLOG/HYDROPLOT system [1] which is nearing obsolescence. SDS III [2] [3] is designed to automate much of the work which is presently done manually, to improve accuracy, to significantly reduce the time from survey to printed chart, and to provide a high degree of reliability and maintainability. Interactive color graphics will be used to aid in surveying and field verification of data. The system will use commercially available, general purpose computer hardware and operating system software. Hydrographic applications software which incorporates the knowledge, experience, policies, and procedures of NOS is being developed jointly by the government and a software engineering contractor. An operational capabilities demonstration is planned for 1987, and it is anticipated that these systems will be used beyond the turn of the millenium.

Two types of systems, both based on Perkin-Elmer 32-bit computers, are being purchased : the Data Acquisition System (DAS) and the Data Processing System (DPS). Installed on survey ships and 9-meter survey launches, the DAS processes and logs data and produces real-time graphic displays to supply steering guidance to the helmsman and to permit monitoring of the progress and quality of the survey. Industry standard hardware interfaces are provided for current and future electronic positioning and depth sounding equipment. The DPS, located aboard the survey ships and at two Marine Centers, processes data collected by the DAS, provides graphic displays and hardcopy plots to help plan, monitor, and modify the survey, and allows the survey officer in the field to make decisions about the significance, validity, and quality of the collected data.

APPROACH

Positioning for SDS III will be handled by a unified 'integrated navigation' approach which can simultaneously utilize multiple lines of position (LOPs) from a selection of short, medium, and long-range positioning systems, as well as from auxiliary heading and speed sensors. Two LOPs are necessary and sufficient for positioning. More, although not required, are potentially beneficial. Fewer than two LOPs, augmented by speed and heading, can be used for limited periods of time before either increasing the number of LOPs to a minimum of two or terminating operations until at least two reliable LOPs become available. Random noise components are reduced and biases recognized by applying advanced filtering (DAS or DPS) or smoothing (DPS) techniques to restrict position solutions according to selected limitations on vessel dynamics. (The term filtering refers to the

use only of past data while smoothing implies the use of both past and future data.) The use of these procedures, which are much more ambitious and calculation-intensive than those in current practice, is made possible by the dramatic increase in available computer power.

LOPs can be derived from any combination of ranges, range (phase) differences, sextant angles, azimuths, and latitude/longitude estimates (i.e., processed GPS data) from sensors such as Falcon 484, Miniranger III, Del Norte R03C and 520, Argo, Raydist, Hydrotrac, Northstar LORAN-C, and Texas Instruments 4100 GPS receiver. Auxiliary data can be derived from gyro, digital compass, Ametek Doppler speed log, and engine RPM and propeller pitch pickoffs. Theodolite angles (azimuths) and sextant angles will initially be entered manually by the operator, although provisions for automated range/azimuth sensors such as «Polar Fix» may be included in the future.

A block diagram of the basic positioning system functions is seen in figure 1. The DAS and DPS versions are very similar, except that the DAS cannot perform smoothing, the DPS need not provide parameters for navigation com-

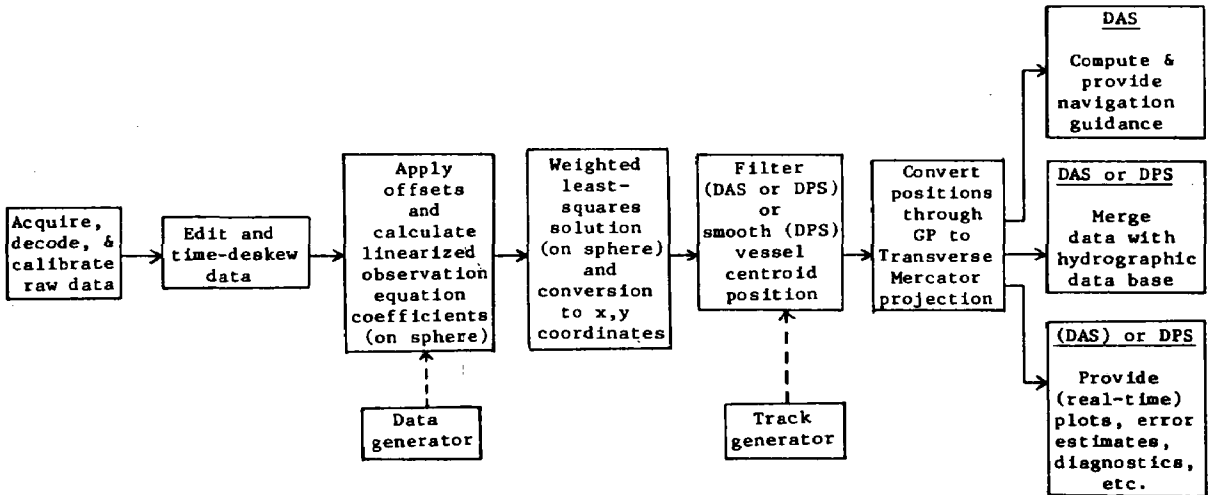


FIG. 1.— Block diagram of basic positioning system functions

mands, and sparse data, which is entered manually in delayed 'real time', is handled differently. All available LOPs and auxiliary data will be calibrated, edited to remove obvious blunders and data with unacceptable signal strengths, and time shifted (deskewed) to common times of interest. Observation equations derived from the LOP data and adjusted for antenna and observer offsets from the vessel 'centroid' are calculated with very high accuracy and a minimum of computer burden by Gaussian conformal mapping from a geodetic spheroid to a sphere [4]. This ensures that coordinate transformation errors remain much smaller than sensor accuracies regardless of range or latitude. Observation equations on the sphere have been developed for range, range (phase) difference, angle (sextant), azimuth (theodolite), and latitude/longitude (GPS) sensors. The equations are linearized about the previous sensor location as the fiducial. There is

no need to predict the present sensor location as the fiducial because of the one-second update rate and typical 5 m/s vessel speeds.

The linearized LOP observation equations are solved via a standard, weighted least-squares [5] adjustment procedure on the sphere to provide an 'unconstrained' vessel centroid position which is converted to a local easting/northing coordinate system. Provision is included for iteration of the solution should the LOPs be tightly curved, as well as for special cases such as inclined-plane sextant angles and situations of severe vertical displacements at short ranges. Weights for each sensor type are primarily fixed in advance but may be adjusted over limited ranges in response to actual conditions. The acknowledged vulnerability to error of least-squares solutions if a poor quality LOP is included, as noted by WEEKS [6], is mitigated here by the following dynamically constrained filter operation which recognizes sudden biases and deactivates offending LOPs, and by the editor which recognizes and deactivates LOPs with excessive random noise.

The unconstrained solution is passed to special-purpose filter or smoother algorithms for reducing the random noise component, recognizing and limiting response to suddenly biased LOPs, and, with the use of speed and heading data, providing the ability to operate for limited periods of time when some or all of the primary positioning data has been temporarily lost. This includes cases of diverse data rates, such as range-azimuth operation, where the manual data is very sparse compared to the electronic data. Partitioning the unconstrained least-squares LOP solution from the following dynamically constrained filtering or smoothing operations permits the above solution iteration prior to filtering and saves computer time through the multiplication of smaller matrices and the implementation of a more efficient smoothing procedure. The filtered position is merged with the hydrographic data base, used to produce plots, error estimates, and diagnostics, and processed in the DAS to provide navigational guidance outputs to the pilot.

A filter can be described as an operation which estimates values of desired output quantities (the state vector) and their uncertainties (the covariance matrix) from a set of noisy measurements of related quantities (the observations vector) based on a model which relates the input and output parameters, the relative accuracy or importance (the weight matrix) of the observations, and the pertinent noise factors. «Measurement» noise is that noise associated with the sensors, while «process» noise represents actual variations in the true states of the vessel compared to an idealized model such as a straight line or a mathematical curve. The object of a filtering operation is to reduce the errors in the state estimates caused by measurement noise on the observations without significantly altering the actual process (vessel motions induced by winds, waves, currents, and steering).

The real-time DAS filter is an augmented version of one developed by HOUTENBOS [7] who combined it with the least-squares solution. His formalism, properly described as a Bayes filter with iterative differential correction, is of the type described by MORRISON [8], but has been extended to include process noise. The outputs or state variables which are estimated are x, y positions and speeds, and speed and heading offsets (between direct sensor measurements and track over the bottom). The Houtenbos approach is unique in that vessel dynamics are limited by invoking a priori pseudo-observations and constraints based on elementary

equations of motion. It includes statistical limitations on vessel accelerations as well as incorporating vessel heading inputs and constraints, when available. We have chosen to add speed inputs and constraints in a similar fashion as well. This provides the added valuable benefit of being able to utilize the heading and speed data for dynamically constrained 'dead reckoning' when primary positioning data is lost for short periods of time. The Bayes formulation is preferred over the mathematically equivalent Kalman formalism because the time-consuming matrix inversions are done in the smaller state space rather than in the larger observations (and pseudo-observations) space.

This improved filter reduces the standard deviation of the measurement noise and provides the opportunity to optimally utilize heading and speed data for positioning. Water speed offset and heading offset are continuously estimated and updated in real time by comparisons with positions derived from the LOPs, when they are functioning. Consequently, even such indirect 'speed' measures as engine RPM (and propeller pitch) can be used successfully with only a rough initial calibration. The speed offset is set to zero if a Doppler speed (over the bottom) sensor is used. The net effect of the environment (winds, waves, and currents) on vessel motion in a given region, relative to measured speed and heading, is termed 'current'. A running estimate of the current is continuously determined from the offsets, and at signal outages the last value is used to augment dead reckoning. The quality of the fix, as determined in part by the covariance matrix and the closure of the LOPs relative to their standard deviations, is monitored and reported in real time in the DAS. This is one of the means by which biases can be detected. Warning flags are set when unacceptable errors accrue.

Further reduction of random noise can be achieved by additional processing (smoothing) off-line in the DPS, because data is then available both before and after each time of interest. A smoother described by Houtenbos requires extremely large matrix inversions and is deemed impractical. An alternate approach, patterned after Mayne [9], has been developed in which the filtered results from the DAS are combined in a weighted least-squares manner with results obtained by running a predictor (similar in nature to the filter) backwards over the 'future' data from the unconstrained least-squares LOP solution. In this way, all available data are used without redundancy. Under typical conditions, the smoother will reduce the random noise component by roughly 30-50 percent over the filtered result. Because processing time for the smoother is roughly double that for the filter alone, the smoother will be optional and invoked only when required, based on accuracy considerations.

To reiterate, the key features are :

- optimal use of overdetermined situations via multiple and hybrid LOPs;
- capability to handle range, range difference, azimuth, sextant angle, and latitude/longitude;
- sensor data 'deskewed' to fixed, common times;
- use of auxiliary inputs such as speed, heading, and signal strength;
- data editing to suppress blunders or 'fliers';
- geodetic calculations one or more orders of magnitude more accurate than data;

- filtering and smoothing with robust algorithms for random noise reduction and bias recognition;
- position solutions statistically limited by permitted vessel dynamics considerations;
- real-time speed and heading offset and 'current' calculation;
- ability to use dead reckoning during data outages of limited length, even during maneuvers, with low error;
- real-time digital displays for helmsman and hydrographer including real-time error estimates and diagnostic messages to the hydrographer.

DESIGN

This section is devoted to a detailed description of the design details for each major process identified in the figure 1 data flow diagram.

Calibration

The raw sensor data is first decoded and calibrated. Positioning system hardware bias errors are calculated by comparing system outputs with standards such as fixed locations known to third-order accuracy or redundant sextant fixes. Calibrators for the sensor data are then derived from the differences through the application of geodetic inverse calculations. Calibrated sensor data will be referred to as measurement data.

Editor/Deskewer

Individual positioning control stations which have previously been established along the shoreline are first denoted as 'used' or 'unused', depending on their status and proximity to the day's survey area. 'Used' measurement data from electronic systems, which arrives at a fixed rate (e.g. 1/second), is inspected for gross blunders or 'fliers' by comparison with an expected value generated by a 'fixed memory' [8] [10] predictor. In effect, the expected value of each arriving data point for each LOP is predicted by extrapolating prior data for that LOP with a linear, least-squares fit over a running 'boxcar' of some fixed length (e.g. the 12-15 preceding values). In practice, this procedure is carried out very efficiently through the use of an expression of the form :

$$p_{L(i+1)} = \sum_{j=0}^N w_{M,N(i)} d(j) \quad (1)$$

where :

- $i+1$ represents the arrival time of the current data value being tested;
- $p_{L(i+1)}$ is the predicted measurement value for the L^{th} LOP at time $i+1$;
- N is the length of the $N+1$ point 'boxcar' buffer whose last point is at time i ;
- $d(j)$ is the j^{th} data value in the buffer; and
- $w_{M,N(i)}$ is a set of $N+1$ predetermined weights which depend on the length of the predictor, N , and the time of the prediction, M , from the same origin as N .

For a one-step prediction past a length, N ,

$$w_{N+1,N(i)} = 2(3j-N)/N(N+1). \quad (2)$$

A priori estimates of the standard deviation of the random measurement error, σ_{ap} , (in meters) are established for each LOP. Estimates not already in units of meters (e.g. in a phase-difference system) are converted to meters using the approximate position and knowledge of the local field gradients of the LOP patterns. If the difference, $\Delta(i+1)$, between the most recently measured value and its corresponding predicted value, is less than $n\sigma_{ap}$, where n is a parameter which may fall roughly in the range (5-20) depending on the sensor type and signal conditions, then the measured value is accepted, passed through unaltered, and entered as the last point in the buffer in preparation for the arrival of the next measurement. If, however, $\Delta(i+1) > n\sigma_{ap}$, then the measured data value is rejected and replaced in the data stream and buffer by the predicted value.

At initialization, electronic data passes through, unedited, directly to the deskewer until a minimum buffer length (e.g. 5 points) is filled with 'good' data whose standard deviation about the mean does not exceed some (other) threshold value multiplier of the a priori standard deviation. In this way, the LOP remains inactive until fliers are, as they must be, excluded from the initial data buffer. The editor then becomes operational and increases in length with the addition of each data value until its customary maximum size is reached. With those LOPs for which signal strength is a reported parameter, predicted values also replace measured values whose signal strengths do not exceed some preset threshold.

It must be emphasized that the editor is not a filter. Most of the data is expected to pass through unaltered; only gross blunders, many a priori standard deviations from the predicted LOP value, or data with unacceptable signal strengths are replaced. The editor may be purposefully bypassed if severe maneuvers, which might cause good data to be rejected, are either preplanned or sensed from auxiliary data.

The number of consecutive replacements for each LOP, for whatever reason, is tallied and limited, because a set of predictions longer than the predictor buffer will diverge from the true process mean as a function of time. In such a case good data values, upon their resumption, might not be recognized as such, thereby causing a 'runaway' condition. Thus, if the number of consecutive predictions exceeds a certain threshold, the offending LOP is declared 'inactive' until such time as the data once more becomes reliable. The hydrographer is informed by a warning message and a buzzer. Inactive LOPs are continuously monitored and are allowed to return to 'active' state when a reduced length predictor data buffer (say, e.g., 5 values) has been refilled with 'good' data. That fact is announced to the hydrographer via a message and a bell. As at initialization, the data buffer of the predictor then grows one step at a time to its full length.

The standard deviation about the mean of the differences, $\Delta(i+1)$, over a selected data set length, is monitored and compared against three criteria. Firstly, if the standard deviation exceeds a certain absolute value in meters (based on the scale of the survey and the intersection angles with other LOPs) for a given conse-

cutive number of points, a red warning flag is sent with a message to indicate that the required survey accuracy is in danger of being compromised. Secondly, if the calculated standard deviation differs from the a priori standard deviation for that LOP by more than some selected multiplicative factor for a fixed number of consecutive points, a yellow warning flag and message is sent to indicate to the hydrographer that the a priori value, which is used in the filter to determine the relative weights of the LOPs, may no longer be appropriate.

Thirdly, the standard deviations from all active LOPs are monitored to insure that the information from each of multiple LOPs is beneficial rather than detrimental to the overall positioning accuracy. To accomplish this, the largest standard deviation is compared with the mean of the two smallest standard deviations, and if the ratio exceeds a certain threshold for a given consecutive number of times, a yellow warning flag and message are relayed to the hydrographer that the LOP with the largest standard deviation could be declared inactive. With four or more LOPs, this process is repeated with the new largest standard deviation until all remaining LOPs are judged by the above criterion to be improving the solution or until only the best two are left. All three of these tests continue to be applied to LOPs which become inactive in this manner so that, when conditions improve, green flags and messages are sent to the hydrographer indicating that the LOP(s) in question can once more be made active, or the a priori values readjusted.

Sparse data from optical LOPs bypasses this editor and is sent directly to a separate, similar editor (fixed-memory predictor) of shorter length, whose operation is tailored to the much lower data rates. As with electronic data, performance thresholds are set, and warning messages are sent to the hydrographer regarding poor data quality.

Data from the various electronic LOPs, even for a given positioning system, are not measured simultaneously, but at different times according to the design constraints of the hardware. The crossing of two such 'skewed' LOPs does not, therefore, indicate accurately the location of the vessel at any particular time. In order to remedy this error, all electronic LOP data can be 'deskewed' by extrapolating it to selected common times at which the LOP intersections will provide an accurate estimate of the actual location of the vessel at that time. A fixed-memory predictor algorithm of the type used for editing data blunders can be used to predict LOP values at any desired reference time. The predictor weights required in equation (1) for the more general case of a prediction time increment Δt , past the time of the last measured data value, can be expressed as

$$w_{\Delta t, N(i)} = 2[N(3j-N+1)+3(\Delta t)(2j-N)]/N(N+1)(N+2). \quad (3)$$

Note that for $\Delta t = 1$, this expression reduces to equation (2). Because a great deal of historical data is neither necessary nor desirable for this purpose, the length of the deskewer will be about 5 points, shorter than that for the editor, to increase responsiveness.

Reference times for the deskewing can be either calculated and set internally from examination of the LOP temporal arrival pattern or dictated by optional operator input. Reference times are indicated to the hydrographer by a flash-

ing indicator which can be used for sending voice 'mark' commands for on-board sextant fixes or via radio to shore-based theodolite operators measuring azimuths. Because of its imprecise temporal nature and low rate (typically one measurement per minute), sparse optical data is not deskewed in the same manner. It has been noted in field tests [11] that measurement errors equivalent to a time delay of about one second between the 'mark' and the measurement are typical for shore-based azimuths. This so-called 'Waltz delay', or one of appropriate magnitude, can be applied to the recorded data times to automatically compensate for this effect, thus providing improved accuracy in azimuth data. A similar procedure can be used for sextant angle measurements.

Heading data from magnetic or gyro compasses, which is hard-wired to the DAS, is deemed to require no editing or deskewing. Loss of data due to hardware malfunction would be handled either by operation without heading data or, preferably, by stopping for repairs.

The primary source of speed data, certainly for the launches, will be engine RPM converted to water speed via a rough, steady-state calibration relationship. This calibration, which is known to vary over short and long term with a number of factors such as direction of the vessel with respect to the prevailing sea and wind conditions, hull cleaning and engine maintenance, etc., acts only as an initialization tool. Thereafter, the offsets between RPM-derived water speed and ground speed, determined from successive vessel positions, are updated continuously by the filter. For large vessels with variable pitch propellers, the pitch must also be monitored and included in the water speed estimation. The engine RPM and pitch are also hard-wired and not believed to require editing or deskewing. It is expected that the RPM data will be low-pass filtered (damped) in such a way as to better match vessel speed to RPM during highly dynamic conditions such as acceleration from a standing start.

The only other currently envisioned potential speed sensor is a two-axis sonar Doppler speed log which may measure bottom speed or water speed depending on the depth of the water. These have various modes, internal filters, data loss indicators, and processing procedures of their own which would further complicate the SDS III algorithms. Field experience with these sensors on NOS vessels indicates that the outputs are frequently noisy and unreliable and would have to be carefully edited and annotated. It is felt that such instruments might best be included in the SDS III sensor suite at a later time.

Observation Equations

Measurement data (ranges, range or phase differences, azimuths, angles, and latitude-longitude) must be converted into lines of position (LOPs) on some reference surface. The reference surface chosen to maintain a high degree of accuracy under all circumstances is a sphere with parameters specially defined to best fit the geodetic spheroid in the survey region of interest. The given positioning system control locations are transferred to the sphere by Gaussian conformal mapping [4], [12]. On the sphere, LOPs for each active sensor type are constructed from the measurement data using equations derived from spherical geometry [13]. Ranges and range (phase) differences can be corrected for estimated speed-of-light variations and geoid heights [5]. All signals except azimuths are corrected vertically for predicted tides.

The LOPs are linearized in the region of the estimated or 'fiducial' sensor antenna or observer position. Horizontal offsets on the vessel from the antenna or observer station to the 'centroid' of the vessel are automatically taken into consideration by use of the heading data. Upon initialization, this fiducial is entered by the hydrographer and is merely a rough estimate which should be good to a few hundred meters or roughly ten percent of the distance to the nearest control station, whichever is smaller. Once underway, the fiducial is generally the previous position estimate out of the filter, corrected to the antenna or observer location, or, if the LOPs are highly curved in that region or if there are significant height differences, it may be the offset-corrected, unconstrained least-squares solution fed back for iteration. With these design considerations, even difficult geometries such as inclined-plane sextant angles, ranges or range differences with large vertical height differences, and large vessels (whose antenna or observer positions may be far from the vessel centroid) operating near control stations are handled automatically with high accuracy.

The general form of all observation equations is :

$$\frac{\delta F}{\delta \Phi} \Delta \Phi + \frac{\delta F}{\delta \Lambda} \Delta \Lambda = F_m - F(\Phi_0, \Lambda_0), \quad (4)$$

where $\Delta \Phi$ and $\Delta \Lambda$ are corrections to the nominal or fiducial spherical latitude and longitude, respectively, and F is the computed value of the measurement, or some function thereof, when the centroid of the vessel has spherical latitude and longitude, (Φ_0, Λ_0) . For example, the range to a station is used per se for F , but it may be computationally more convenient in other cases to use the cosine of a sextant angle or the product of a measured azimuth angle with the approximate range to the vessel from the observing station. The partial derivatives on the left-hand side are computed at the point (Φ_0, Λ_0) , and F_m represents either the measurement itself or the function of the measurement, as represented by the function F . The weight of the observation equation for the following unconstrained least-squares calculation is taken to be the inverse square of the estimated error in F_m . Specific equations for each sensor are listed in Ref. [13].

The linearized observation equation coefficients in spherical coordinates are sent, under most circumstances, to the unconstrained least-squares algorithm for position estimation. For several special cases, such as only one active LOP or sparse data with two or more non-sparse LOPs, the equations are converted to a local easting-northing coordinate system and sent directly to the filter, bypassing the least-squares calculation.

Unconstrained Least-Squares Solution

In the original HOUTENBOS [7] formalism, the weighted least-squares and the dynamically constrained filtering operations were combined into a single procedure. These operations are basically independent, as can be noted by the way in which the matrices partition. We have chosen to separate the operations and apply them serially in order to improve computing efficiency (due to the multiplication of smaller matrices and to a more efficient smoother algorithm) and to permit iteration, if necessary, prior to filtering.

The weighted least-squares position solution is termed 'unconstrained' because it has been separated from the dynamic constraints (which limit changes in the position of the vessel) which are now imposed during the following filtering step. It is a standard mathematical procedure which can be referenced in any pertinent text and need not be further discussed except to note that this implementation is performed differentially, i.e., the least-squares results are increments, $\Delta\Phi$ and $\Delta\Lambda$, from the previous solution. The weights are the inverse squares of a priori estimates of the LOP errors in meters and are constant (unless altered manually) except for the case of range or phase difference where the difference error is assumed to be constant, and the error in meters is calculated as a function of the location in the LOP pattern through use of the scalar field gradients. The least-squares procedure also yields a solution covariance matrix which is scaled to x, y units (as below) and used to supply weights for the unconstrained solution components in the filter algorithm that follows. The information in the covariance matrix can be converted to an error ellipse [12] which may be displayed to the hydrographer.

This block of the computer code also : 1) calculates LOP intersection angles to assess error magnification due to geometric dilution of position, 2) calculates the scalar field gradients needed in the editor for conversion of phase differences to units of meters, 3) controls iteration of the solution as a new fiducial back to the observation equation calculation for cases of difficult geometry, and 4) converts the incremental position solution from spherical latitude and longitude ($\Delta\Phi$, $\Delta\Lambda$) at time step, i , to local easting and northing increments (Δx , Δy) via the basic equation

$$\Delta y(i) = R \Delta\Phi(i) \quad (5)$$

$$\text{and } \Delta x(i) = R \cos\Phi(i) \Delta\Lambda(i). \quad (6)$$

These increments are added to the fiducial x, y values to yield the updated x, y position.

Dynamically Constrained Filter/Smother

The primary positioning sensor data for multiple LOPs, reduced via a standard, weighted least-squares algorithm, provides easting and northing (x,y) position estimates based solely on the noisy sensor data and unconstrained by limitations on vessel dynamics. The HOUTENBOS [7] method of applying vessel dynamics constraints in the generation of filtered paths involves the use of elementary equations of straight-line motion as constraints along with the actual input data. The basic assumption or pseudo-observation is that while the helmsman is attempting to maintain a straight-line course at a speed 'V', the mean acceleration is zero — otherwise the course would be curved or the speed varying. The actual accelerations experienced by the vessel (the process noise) are modeled statistically as isotropic with an estimated standard deviation, σ_a , about the zero mean. This value is used in determining the weighting factors for the pseudo-observations. Turns whose centripetal acceleration (V^2/r) is no greater than roughly σ_a are also accommodated by this model.

The so-called 'state vector' is composed of the quantities being estimated; i.e., it is the answer. In this case, the state vector is defined as

$$Y = (x, y, u, v, b_h, b_s)^T.$$

where x and y are the easting and northing components of the filtered position solution, u and v are their respective speeds, and b_h and b_s are the heading and speed offsets. The transpose notation (T) is invoked simply to write this column vector horizontally on the page to save space. Heading and speed are modeled via observation equations including offsets between measured quantities and the values over the bottom derived from the actual path or 'track made good' as determined from the speed components of the filtered solution. The model assumptions or pseudo-observations in the constraint equations are that the mean rates of change of the heading and speed offsets are zero with standard deviations, σ_h and σ_s which are used in determining weighting factors. Based on a straightforward analysis [12] the latter values have been coupled to σ_a via the relationships $\sigma_h = \sigma_a/2v$ and $\sigma_s = \sigma_a/2$. In the weight matrix these values are multiplied by heading and speed 'variance multipliers' which provide control over the coupling to σ_a in terms which can be related to perceived vessel motions. The heading variance multiplier is calculated from a 'turning factor' which expresses the estimated fraction of cross-track vessel movement caused by heading changes. The speed variance multiplier is calculated from a 'correlation factor' which relates the estimated response of the speed sensor (i.e., engine RPM) to in-track movements caused by the wave field.

For this filter [12] [13], based on a linearized, iterative differential correction approach, the object is to determine the changes in the state vector, Y , caused by changes in the measurements or observations change vector, ΔX . The 'system' coefficient matrix, A , for the linearized observation equations is defined through the relationship $dX = AdY$, where the elements of A are $a_{mn} = \delta X_m / \delta Y_n$. In the Houtenbos approach, the approximate or estimated state values to be updated at the present iteration are not predicted formally, but rather, because of the frequent update rate and low speeds, simply set equal to the solution values at the previous iteration. The measurements change vector (observed minus estimated), composed of four true observations and six pseudo-operations, is constructed as follows :

$$\Delta X(k) = [x_{u(i)} - x(i-1), y_{u(i)} - y(i-1), h_{m(i)} - h(i-1), s_{m(i)} - s(i-1), u(i-1)t, v(i-1)t, 0, 0, 0, 0]^T.$$

where i indicates the iteration time step number, h is heading, s is speed, the u subscripts denote the unconstrained solution, the m subscripts indicate measurements, unsubscripted quantities at step $i-1$ are the previously estimated quantities, and t is the time between data updates. The formal, dynamically constrained least-squares solution or state vector update is

$$Y(i) = Y(i-1) + H(i) A(i)^T W(i) \Delta X(i) \quad (7)$$

where W is the combined measurements and constraints weight matrix, and H is the state covariance matrix which is calculated as $(A^T W A)^{-1}$. The constraints weight matrix is updated at each step by adding the state covariance matrix from the previous step to the covariance matrix of the acceleration effects.

There are a number of possible combinations of numbers and types (electronic and optical) of LOPs which are handled as special cases. The most important of these are the cases of either zero or only one active LOP. Because no unconstrained position can be determined directly from measurement data alone under these circumstances, the speed and heading offsets can no longer be directly determined and must be removed from the state vector, along with the corresponding rows and columns from the system, weight, and covariance matrices. If this is not done, the system of equations is underdetermined, and errors are partitioned, as rapidly as dynamic constraints permit, into both the offsets and the solution.

In order to compensate for this loss of real-time information, the effects of wind, waves, currents, etc., on the vessel, as they affect the difference between measured speed and heading and the actual track, are reduced to a single net mean vector termed 'current'. Its x and y components are constantly updated from the speed and heading offsets, while at least two electronic LOPs are active, by averaging over a given period (30-120 sec) of preceding data, i.e., 'running boxcar' means. The standard deviations about the mean 'current' components are used in the weight matrix, W. It can be noted that the lateral accelerations which cause the true path to differ from the planned path lead to a speed bias, because they increase the path length but are not sensed by the speed measurement. The value of this bias is proportional to $(\sigma_a/V)^2$. This speed bias must not be allowed to become part of the speed offset but rather included as a separate term in the speed equations so that the speed offset which results from the net 'current' can be used to calculate an unbiased estimator of the 'current'.

When the number of active electronic LOPs drops below two, the value of the 'current' remains fixed at its last estimated value for application to the dead reckoning calculations in the filter involving measured speed and heading. For fewer than two active electronic LOPs, the measurement change vector is rewritten in terms of the current components, c_x and c_y . Whenever a two LOP fix is obtained, even for just one sparse point, the value of the 'current' is modified to reflect the observed difference between the new solution and the filtered prediction.

If the sensor complement drops to zero active electronic LOPs, the appropriate matrix elements are zeroed out, and the algorithm will continue to provide positioning and navigation outputs through the use of speed and heading data, the previously estimated net 'current', and the dynamic constraints imposed on vessel motions. The algorithm, which is the result of examining the performance of a number of diverse formulations, has been designed specifically to function well even during maneuvers such as U-turns. Errors grow with time, and operation cannot continue indefinitely without accruing an unacceptable positioning bias. Estimates of the expected positioning error derived from the measured uncertainty in the 'current' estimate and the LOP outage time are presented to the hydrographer to aid his judgement as to when operations must be terminated until LOPs can be reactivated.

For the one active electronic LOP case, since no unconstrained solution is possible, the measurements and observation equation coefficients are sent directly to the measurements change vector and system matrix, respectively, in the filter. The processing procedure depends on whether the one-LOP case was arrived at from two LOPs or from zero LOPs. For the former case, the one-LOP data is used in the solution. This mode of operation, which might be termed 'augmented dead reckoning', provides superior performance over a longer period of time than dead reckoning with no incoming LOP data, but as with that case, errors grow with time. It is designed chiefly for the range-azimuth mode, but it also permits operational capability while waiting for lost LOPs to be reactivated. As with the zero LOP case, positioning error estimates are continuously updated and provided to the hydrographer. For the zero-to-one LOP case, the dead reckoning position has already accrued potentially significant errors, and the addition of one LOP would, in general, not be particularly useful. Indeed, it could, for example, cause an apparent reversal in the ship's track. For this reason, the one-LOP procedure, when coming from zero LOPs, will be to simply report the residual (from the dead reckoning position to the LOP) to the hydrographer and to continue dead reckoning.

Sparse optical data must be handled differently from electronic data because of the highly diverse data rates, and because the sparse data is entered manually into the computer some time (e.g. 10-30 seconds) after it is measured. In the DAS, estimated positions are buffered for that duration to permit them to be adjusted, before storage, for the time period between the measurement of the optical data and its entry into the computer. At the times of the optical data measurements, the dynamically constrained, augmented dead-reckoning positions estimated at a 1 second rate will not agree precisely with the actual fix obtained including the optical measurements, but this will not be known until the data is entered manually and processed. At that point, the position solution at the time of the measurement is reset to the unconstrained solution. The estimate of the 'current' is updated, new position estimates based thereon are calculated beginning at the time of the actual fix, and the new position estimates replace the earlier temporary values up to the actual clock time. In the DPS, the corrected values of 'current' necessary for closure at the sparse, unconstrained solutions can be precalculated (by looking ahead in the data) and applied immediately at the actual measurement time such that the augmented dead-reckoning solution will merge with the sparse, two LOP fix.

The data processing procedures vary significantly depending on the number of active sparse and electronic LOPs. These are summarized in Table 1.

The values of σ_a , σ_h , σ_s , and the variance multipliers in the constraints weight matrix act as tuning parameters on the filter which regulate the rate at which changes in indicated position, velocity, and heading and speed offsets can occur. Small values of these parameters produce very smooth tracks which approximate the true path along straight lines but which cannot react without overshoot to rapid changes in course and speed. Large values of σ_a yield the ability to follow maneuvers but do not provide sufficient reduction of random measurement noise. Depending on the size of the vessel, the sea state, and the positioning system error, there generally exists a compromise value which provides adequate random noise reduction on lines without causing unacceptable biases

on turns. This is true because in typical hydrographic operations, positioning accuracy on turns is not critical, and the only real requirements are to maintain lane count in hyperbolic systems and to shed biases and get back on line quickly after the maneuver. The availability of heading data greatly aids the filter in maintaining a low bias condition during turns.

Table 1

Processing Procedures

Key : A = No. of active, sparse LOPs,
 B = No. of active, high-rate (1/sec) LOPs, and
 C = Total No. of LOPs.

A	B	C	Procedures
0	0	0	Dead reckoning (DR) for limited time; try to reactivate one or more LOPs; if not, stop.
0	1	1	From 0/≥2/≥2 case, augmented dead reckoning for (longer) limited time; try to reactivate another LOP; if not, terminate. From 0/0/0 case, continue DR and report residual from DR position to LOP.
1	0	1	Dead reckoning; sparse LOP not used in solution but yields residual to DR position for limited time; try to reactivate another LOP; if not, stop.
≥1	1	≥2	Augmented DR until arrival of sparse LOP(s); then calculate unconstrained solution; calculate difference between solutions for use in modifying 'current'; reset solution to unconstrained solution; continue.
0	≥2	≥2	Normal operation; if from 0/1/1 case, reset position to unconstrained solution, recalculate current, and continue.
≥2	0	≥2	DR between coincident sparse data; upon arrival, use sparse data unconstrained solution to replace DR solution and recalculate 'current' from difference between solutions.
1	≥2	≥3	Normal operation; sparse LOP not used in solution but, at arrival of sparse data, report residual from filter solution to sparse LOP.
≥2	≥2	≥4	Normal operation; sparse fixes not used in solution, but residuals from filter output to unconstrained sparse position solution reported.

The primary logistical factor which must be considered in designing a survey is to allow sufficient time on line after major accelerations (turns and speed changes) for the biases to damp out before highly accurate positions are

required. The greater the noise reduction (small σ 's) the greater the time needed for equilibration. Typical times for launches might be 10-30 seconds, while for ships it could take as long as 40-60 seconds for very small σ_a . If it were desired to further reduce the biases accrued during planned turns (at the cost of less measurement noise reduction) for some specific application such as lines ending in shoal waters, the values of the tuning parameters could be increased in real time according to navigation requirements. As will be seen, this is unnecessary in most cases because, with speed and heading data available, adequate noise reduction can be achieved without significant added bias.

For output, the filtered positions in x , y are converted back to spherical representation using equations analogous to equations (5) and (6) :

$$\begin{aligned} \Phi_{f(i)} &= \Phi(i) + [y_{f(i)} - y(i)]/R & (8) \\ \text{and} \quad \Lambda_{f(i)} &= \Lambda(i) + [x_{f(i)} - x(i)]/R \cos \Phi_{f(i)} & (9) \end{aligned}$$

where 'f' subscripts denote filtered values, and non-subscripted variables are the unconstrained solutions.

Coordinate Conversion

The output of the filter is in spherical latitude and longitude. It is desired that the various output products, such as plots and data bases, be in a standardized, planar coordinate system or map projection such as Transverse Mercator (TM) [5]. The filter output could be converted directly from the sphere to TM, but this could lead to accuracy problems because the various projections are not associative. In other words, going directly from geodetic position (GP) coordinates (latitude/longitude on a standard reference spheroid) to TM will produce a slightly different result than going from GP to the sphere and from there to TM. Because SDS III is concerned with historical and ancillary products which may have been externally converted directly from GP to TM, the positioning results should, in order to maintain the highest possible accuracy and consistency, be handled similarly. The filtered positions expressed in spherical coordinates are thus first transformed back to GP and then to TM using standardized transformation formulæ.

Simulator

In order to quantify performance against known quantities and prior to the availability of actual field data, the algorithms have been exercised firstly via a 'track generator' which simulates the output of the preceding weighted least-squares solution converted to x - y coordinates to exercise the filter/smoothing alone, and secondly by a 'data generator' which simulates noisy LOP data for testing the observation equations and unconstrained least-squares code as well. Initially, a 'planned' ideal path selected from a menu containing a straight line, a 90-degree turn, a U-turn, a S-turn, or a racetrack is constructed. The vessel travels at a selectable constant speed through the water. Process noise representing actual wind/wave-induced vessel track and speed deviations for random accelerations of selectable magnitude and five-second duration is calculated in along-track and cross-

track components. The along-track magnitude is permitted to differ from the cross-track magnitude by a selectable factor to permit simulation of various wave fields and attack angles. The resulting deviations are interpolated to one-second intervals, converted to x and y deviations with the use of the heading information, summed, and applied incrementally to the x and y components of the planned path to produce the 'true' path. The five-second duration, selected to be representative of the yaw rate of a survey launch under moderate sea conditions, provides desired cross-track deviations in the 5-10 m range for a cross-track acceleration of 0.5 m/s^2 .

Instantaneous vessel heading values are calculated from the true path by invoking a selectable coupling factor appropriate for the size of the simulated vessel and the update period. Two types of simulated heading sensor errors are generated. The first is random one-second deviations of selectable magnitude corresponding to roll and pitch effects. The second is compass bias errors (due to its damping and subsequent delayed dynamic response) with magnitudes initially equal to the five-second process noise-induced course changes times a fractional multiplier called the compass damping factor. These values are linearly damped to zero in five seconds and interpolated to one-second values. Random speed errors of 5-second duration and variable magnitude, interpolated to one-second values, are applied to the assumed speed through the water to simulate speed measurements. Simulated water currents may be applied to skew the path and yield different speeds over the bottom. For ease of computation, the currents are permitted to distort the planned path rather than the more complex case where headings would have to be modified in order to recreate the undistorted original. This shortcut has no effect on subsequent filter analysis or performance.

For the 'track generator', random measurement errors with a priori standard deviation, σ_m , representing apparent deviations in the vessel track caused by random noise in the positioning sensor systems, are added to the x and y components of the true path to yield the 'measured' path which is the input signal for the filter. The measurement errors consist of two populations: the typical, limiting random noise expected during normal operating conditions, and a 'flier' population of selectable probability and magnitude which can introduce the infrequent but much larger spurious responses not uncommon in some systems. Antenna motion caused by vessel roll and pitch, although not totally random, is considered to be part of the measurement noise magnitude. Provision has been made to simulate data drop-outs by turning off any combination of x, y, heading, or speed inputs halfway through the run.

For the 'data generator', the position increments of the true path are geodetically inverted, according to the appropriate observation equations for each sensor type, to yield the corresponding LOP data values. Appropriate random measurement noise is added to these 'true' LOP data values to produce characteristically noisy data. This is sent to the positioning algorithms, beginning with the calculation of observation equations, to test all following code. If varied observation time were added, the editor/deskewer could be exercised in a similar fashion. At this point, neither time nor a mixture of high-rate and sparse data have been simulated.

PERFORMANCE

One measure of the performance of a filter or smoother is the ratio, R , of the standard deviation of the output about the true path to the standard deviation of the input (measurement) noise. These so-called 'filtering or smoothing ratios' are functions of the ratio, Q , of the assumed process (acceleration) noise to the measurement noise ($Q \equiv \sigma_a t^2 / \sigma_m$). The numerator of Q depends on the size of the vessel and the sea state while the denominator depends on the positioning system. Theoretical performance curves for prediction, filtering, and smoothing with no heading or speed inputs, as reported by Houtenbos for the case of random process noise at the measurement period, are depicted in figure 2. These levels of performance have been confirmed with simulated data inputs conforming to the Houtenbos noise model. If the ratio, Q , is small, the measurement noise dominates, and the filter or smoother will be able to reduce its magnitude. If Q is large (i.e., unity or above), the measurement noise cannot be distinguished from the actual ship motions, and the filter and smoother become ineffective (R approaches unity). Note that for large Q the predictor actually degrades performance as R increases above unity.




FIG. 2.— Noise reduction factors for predictor (P), filter (F), Houtenbos smoother (S_1), and bidirectional smoother (S_2).

The theoretical performance of the hybrid smoother (forward filter, backward predictor) for no heading or speed inputs, calculated from the predictor and filter curves, is indicated as a dashed line. It is a distinct improvement over the filter, particularly at small Q , but reflects the poor performance of the predictor at large Q . Although R for the hybrid smoother is slightly larger than for the theoretical Houtenbos smoother, it is probably no worse than for a practical implementation of the Houtenbos technique.

The availability of heading and speed inputs further lowers the filtering and smoothing ratios by amounts depending on the heading and speed error magnitudes via assumed values of σ_h , σ_s , and the variance multipliers. Heading information reduces cross-track positioning errors, while speed data reduces along-track positioning errors. Heading and speed inputs are also valuable in reducing biases accrued during maneuvers and in providing 'dead reckoning' information when the electronic positioning system signals are lost for short periods of time.

Figure 2 cannot be directly applied to the more general case where the process noise is not random at the measurement update period. For SDS III, the LOP update rate will be once a second, while it is felt that typical vessel yaw motions are more appropriately represented by a roughly five-second period. The actual performance is expected to lie between the now overly optimistic Figure 2 value based on σ_a and a pessimistic value obtained by replacing σ_a with $5\sigma_a$ (the value needed to yield the same total five-second track deviation in five summed one-second pieces).

Figure 3 includes plots of the a) 'planned', b) 'true', and c) 'measured' paths from the track simulator for a straight-line case of 60 measurements at one-second intervals starting at the bottom of the figure with a constant 5 m/s speed. Process noise accelerations are isotropic at 0.5 m/s² RMS and yield maximum off-track deviations of about 10 m. This represents, for example, a 9-m survey launch in moderate seas. Simulated compass lag errors modeled for the inability to respond instantaneously to track direction changes are as large as 15 degrees with an RMS of about 8 degrees. Measurement noise in x and y is rectangularly distributed with a 5-m standard deviation and no fliers.




FIG. 3.— Track simulator output paths.

FIG. 4.— Filter outputs compared to true path for 0.5 m/s² RMS process noise and $\sigma_a = 0.5$ m/s².

In Figure 4, filtered paths for $\sigma_a = 0.5 \text{ m/s}^2$ are compared, without (4a) and with (4c) heading data, to the 'true' path (4b) which was generated with 0.5 m/s^2 process noise. For the latter case, the heading variance multiplier is 0.09, and estimated random heading error is six degrees RMS. Filtering ratios in the cross-track direction for the two examples are 0.48 and 0.32, respectively. Figure 4c exhibits significantly reduced random excursions and, as it should, clearly reflects the character of the process noise contaminated 'true' path. $R = 0.40$ can be achieved for $\sigma_a = 0.1 \text{ m/s}^2$ and no heading data, but typical turns of 0.5 m/s^2 would not be successfully negotiated at this value. Reducing σ_a to 0.1 m/s^2 for the case with heading data again yields $R = 0.32$, the same value obtained at 0.5 m/s^2 . This is because the actual vessel excursions of 0.5 m/s^2 RMS are now being treated improperly as measurement noise and partially filtered out. Thus, for values of σ_a appropriate for maneuvers, the filter reduces the measurement noise by a factor of two without heading data and by a factor of three with heading data, for this particular parameter set. The use of speed data produces similar improvements in R for the in-track direction, although the percentage improvement is typically somewhat less due to the consideration that a larger variance multiplier is deemed appropriate.




FIG. 5.— Filter output for 50-m radius U-turn at 5 m/s.

Figure 5 demonstrates the utility of heading data in improving overall performance in a U-turn situation with a 50-m radius and 0.5 m/s^2 acceleration. Figure 5a depicts the 'measured' path with 5-m RMS errors. Figure 5b superposes the filtered path for $\sigma_a = 0.5 \text{ m/s}^2$ for the case of no heading or speed data over the 'true' path with its 0.5 m/s^2 RMS process noise. This situation is an improvement over Figure 5a, but it retains a large component of the measurement noise character. Reducing σ_a to 0.3 m/s^2 , as seen in Figure 5c, leads to a slightly smoother but clearly biased result with maximum deviations of 9 m and 13 m in x and y. Although these are probably acceptable in the field under most conditions, Figure 5d demonstrates the added gain from heading data. With σ_a remaining at 0.3 m/s^2 , slightly less than the actual turning acceleration and process noise, the filter utilizing heading information has produced a result which nearly removes the measurement noise but clearly retains the character of the process noise — the actual vessel motions!

This filter algorithm has proven to be very stable and robust. Attempts to cause it to 'blow up' on unrealistically extreme maneuvers with mismatched parameters have failed. Although large temporary biases are incurred in such instances, the algorithm continues to function in a reasonable manner and returns as quickly as it can to a satisfactory solution.




FIG. 6.— Measured and filtered paths for U-turn with dead reckoning after loss of positioning data.

In instances where all primary positioning sensor data are lost during operations, the algorithm will use speed and heading inputs and the estimated 'current' to predict position. The modified, zero-LOP algorithm is expected to provide useable results with primary data drops for periods of more than 30-60 seconds. Useable here means that errors do not grow large enough to cause lane drops in phase measurement systems. Figure 6 is an example of a U-turn in which data from a 5-m RMS, 2 LOP positioning system are lost halfway through the turn. The filter, augmented with noisy speed and heading data and a 'current' estimate (zero in this case), successfully completes the turn, and 30 seconds later, at the end of the simulation, the positioning errors from the true path are only 2 m in x and 5 m in y. These errors, which depend on the speed and heading noise and on the uncertainty in the estimated 'current', are considered to be representative of typical operational conditions. The RMS heading errors were 3 degrees, true, and 6 degrees, estimated (to account for a compass damping factor of 0.8); and the speed errors were 0.25 m/s, true, and 0.15 m/s, estimated. Heading and speed variance multipliers were 0.09 and 0.50, respectively. Changing the speed variance multiplier to 0.95 caused no significant change in the final error magnitudes. These results are certainly far more than satisfactory considering the extreme conditions under which they were generated and the performance requirements under such circumstances.

Figure 7 displays the additional increase in accuracy which can be achieved through use of the bidirectional smoother. The parameters were set for less than optimum performance (e.g., no heading or speed data used and σ_a larger than necessary) to better demonstrate the effects visually. For this case with 0.5 m/s^2 process




FIG. 7.— Example performance of the filter and smoother for 5-m measurement noise, 0.5 m/s^2 process noise, and $\sigma_a = 0.5 \text{ m/s}^2$; no heading, speed.

accelerations and an estimated σ_a of 0.5 m/s^2 , the noise reduction ratios, R , for the filter in (x, y) are 0.49 and 0.59, and those for the smoother are 0.24 and 0.47. The theoretical equilibrium values predicted by the model are $R_f = 0.60$ for the filter and $R_s = 0.45$ for the smoother. The differences of the measured values from these theoretical performance predictions reflect the statistics of the particular arbitrarily selected random noise sequences used, and the fact that the model assumes one-second (the measurement rate) uncorrelated noise, while the process noise is really applied on a 5-second basis to more closely resemble physical expectations. The smoother provides an extra increase in measurement accuracy when needed, but at the expense of doubling the processing time over that for the filter alone, because of its bidirectional nature. The visual evidence of performance in Figure 7 is certainly striking. The smoother has reproduced an amazingly faithful rendition of the true path which was heavily contaminated and made almost unrecognizable by measurement noise.

CONCLUSIONS

SDS III will be equipped with an integrated positioning architecture which makes optimal use of information from a wide variety of electronic and manual positioning devices as well as heading and speed sensors. Data from multiple lines of position from a variety of diverse sensors are combined in a least-squares algorithm followed by a sophisticated Bayes filter which utilizes measured speed and heading data and which invokes knowledge of the dynamic limitations of vessel motions. Heading and speed sensor offsets and the effective 'current' are continuously estimated and updated. Random measurement errors are greatly reduced without significant alteration of actual vessel motions. Bias errors caused by poor data are detected and automatically corrected to the greatest extent possible. In the event of total loss of primary positioning data due to null zones or other temporary problems, the algorithm will continue to supply reasonable position estimates for times long enough to provide a good chance of reacquiring input data before the occurrence of a lane loss or the need to stop due to error growth. Navigation guidance to the pilot and accuracy estimates to the hydrographer are provided in real time. Further improvements in accuracy can be achieved in post-processing with the use of a smoothing algorithm. This software system will provide the hydrographer with higher accuracies and more information, flexibility, and convenience than previously achieved with the same positioning hardware.

REFERENCES

- [1] WALLACE, J. : Fifteen Years of Automated Hydrography in the National Ocean Survey, *Proc. of the Fourth Int. Symposium, 'Hydro 82'*, Southampton, England, December 1982, Vol. 3, Paper 1.

- [2] SCHIRO, R.A. : SDS III — The U.S. National Ocean Service's Next Generation Hydrographic Surveying System, *Proceedings of the National Ocean Service Hydrographic Conference, 'Hydro 84'*, Rockville, Md., April 1984, 177-184.
- [3] ENABNIT, D.B. : Shipboard Data System III, *Proceedings of the First Biennial Conference of the Canadian Hydrographic Service and Hydrographers Association*, Halifax, Nova Scotia, April 1985.
- [4] BAKKER, G., de MUNCK, J.C., and STRANG van HEES, G.L. : Radiopositioning for Marine Geodesy, *Printed Course Lecture Notes*, Delft University of Technology, Department of Geodesy, Jan. 14-17, 1985, p. III-33.
- [5] CROSS, P.A. : The Computation of Position at Sea, *The Hydrographic Journal*, No. 20, April 1981, 7-16.
- [6] WEEKS, C.G. : Field Experience with a New Automated Hydrographic Survey System, *Proceedings of the National Ocean Service Hydrographic Conference, 'Hydro 84'*, Rockville, Md., April 1984, 59-68.
- [7] HOUTENBOS, A.P.E.M. : Prediction, Filtering and Smoothing of Offshore Navigation Data, *The Hydrographic Journal*, No. 25, July 1982, 5-16.
- [8] MORRISON, N. : *Introduction to Sequential Smoothing and Prediction*, McGraw-Hill, New York, 1969, p. 424.
- [9] MAYNE, D.Q. : *Automatica*, 4, 1966, p. 73.
- [10] CASEY, M.J. : Improved Positioning Through Range Filtering, *Canadian Hydrographic Service Technical Report*, 1982, 66 pp.
- [11] WALTZ, D.A. : An Error Analysis of Range-Azimuth Positioning, *Proceedings of the National Ocean Service Hydrographic Conference, 'Hydro 84'*, Rockville, Md., April 1984, 93-101.
- [12] THOMAS, R.W.L. : A Study of Techniques to Improve the Performance of the Shipboard Data System III Positioning System, *Contractors Report WASC TR-X46-003(85)*, EG and G Washington Analytical Services Center, Inc., August 1985, 131 pp.
- [13] THOMAS, R.W.L. : Positioning System Development for SDS III — Observation Equation Derivation and Solution, *Contractors Report WASC TR-8E6-0001(86)*, EG and G Washington Analytical Services Center, Inc., January 1986.