ARTICLE NO. 80-4082

Alternative Approach to Multisensor Navigation

Michael J. Cannon* and Lee A. Hollaar†
University of Illinois, Urbana-Champaign, Ill.

The use of adaptive maximum likelihood (Kalman) filters combining different navigation subsystems to reduce or eliminate characteristic errors is well-established. However, they are difficult to implement on low-cost digital microprocessors, work best with highly refined sensor subsystems, such as inertial or Doppler, and are highly sensitive to both the initial estimates for the model and software errors. A different combining method, based on logical rather than numerical techniques, addresses these problems. It can be implemented on one or more microprocessors (with a fixed communications protocol between the individual microprocessors); requires only limited preprocessing of the input data; can have its constituent subsystems changed without major revision to the combiner as new or better navigation techniques, such as the global positioning system (GPS) or LORAN-C, become available; and tolerates a wide variety of sensor, hardware, and software errors. The basic structure of the method and results of simulations of its performance in the presence of errors common to low-cost sensors (such as VOR or basic Omega) are presented, along with a discussion of possible extensions.

Introduction

As interest in area navigation (RNAV) systems increases, it becomes necessary to develop more accurate position determination systems.[1] Even minor errors in the determination of an aircraft's present position can, in some cases, produce large course deviation indications after transformation by the area navigation equations. The most common radionavigation system, VOR, provides good accuracies in terminal areas and along predetermined radials, such as those used for approaches or airways (although sometimes at the expense of other, nondistinguished radials). However, the errors associated with VOR (which are generally less than 1.9 deg but can be over 6 deg) make it unsuitable alone as the position source for RNAV; generally, it is augmented with more accurate distance measuring equipment (DME) with errors on the order of 0.1 n.mi.[2]

Unfortunately, low-altitude DME coverage is spotty, particularly away from high traffic density areas (and these areas of limited DME coverage are where the ability to provide direct courses and instrument approaches to airports without conventional navigation aids is of the most value), preventing the use of high-accuracy DME/DME position determination, and, in some cases, also VOR/DME. Moreover, as the distance from the ground station increases, the accuracy of VOR/DME becomes marginal. At 40 n.mi. from the station, the error can be as large as 2.5 n.mi.; this error is tolerable only because the certification standards for RNAV[2] are based on the calculated errors for VOR/DME.

Even if DME solved all the problems of position determination for RNAV, the higher cost of an airborne DME transceiver (about three times that of a VOR navigation system) makes it a luxury in those general aviation airplanes that would benefit the most from RNAV's potential.

Other navigation systems do not fare well either. Self-contained systems, such as inertial or Doppler radar, have errors that increase with the length of the flight, requiring the use of a radionavigation system to bound the errors, and are too expensive for most aircraft. Omega[3] provides worldwide coverage from only eight stations, but is heavily affected by propagation disturbances such as sunspots (causing sudden ionospheric disturbances, SID's) and the diurnal shift in the ionosphere's height, requiring extensive processing to achieve acceptable performance. LORAN-C, GPS, and other systems have not been deployed sufficiently to provide satisfactory coverage, and, because propagation disturbances are inherent in any radio transmission system, will also suffer from a variety of errors.

Rather than attempt to produce a system that is free of major problems, with its high development and deployment costs and long phase-in period, a more attractive alternative is to process the existing, flawed information from existing systems to eliminate or reduce errors. The simplest form of this processing is filtering, which can remove errors whose frequency differs substantially from the desired information; the most common example is the dampening of a course deviation indicator to prevent rapid oscillations, such as caused by high-frequency noise. Unfortunately, simple fixed filters are not generally effective and produce a variety of undesired side-effects (such as lags in the response to course deviations).

Kalman Filters

With the advent of airborne digital computers, more complex filtering techniques could be employed to separate accurate information from expected errors. The most common method is the recursively updated maximum likelihood, or Kalman, filter.[4] It consists of a mathematical model of both the dynamics of the system and the expected sensor errors. These are usually expressed in terms of a state vector (which describes the current location, velocity, etc., of the aircraft as well as the calculated nature and magnitude of the sensor errors); time-varying vectors describing the system forcing function and actual errors; and time-varying matrices that map the current state and the forcing function onto the next state, and the current state plus sensor error onto the measured sensor values. Matrix techniques for solving simultaneous linear equations can then be employed to solve for the new state and the sensor errors. (A detailed discussion of Kalman filtering theory[5,6] is outside the scope of this article.)

Not only does the Kalman filter remove the modeled errors, but it allows the use of more than one type of sensor, calculating the relative weighting based on each sensor's perceived performance. This can lead to a system that operates well even in the presence of strong sensor errors (such as severe VOR scallops), provided that some other

---

Received Aug. 9, 1979; revision received Dec. 14, 1979.

Index categories: Navigation, Communication, and Traffic Control; Sensor Systems; Computer Technology.

*Presently, Hewlett-Packard Labs., Palo Alto, Calif.

†Associate Professor, Department of Computer Science, University of Utah, Salt Lake City. Member AIAA.

sensor not affected by the error (such as an air data system) is available. For example, Bryson and Bobick[7] described a simple Kalman filter for combining VOR/DME and air data information which shows an improvement factor of about 2.5 over VOR/DME alone.

However, the Kalman filter combining technique suffers from a number of difficulties. It requires the development of a mathematical model, in the form of the various vectors and matrices, that describes the behavior of the system (aircraft state and sensor errors) at all times when navigation is desired. This may not be difficult for steady-state operation and when a number of simplifying assumptions are made (such as Bryson and Bobick's assumption that the airplane is operating at high altitudes, so that constant VOR/DME coverage is available, and at an airspeed of about 500 knots, so that many signal anomalies can be modeled as white noise).

When real-world conditions (especially reduced speeds in terminal areas or for light aircraft or helicopters, and low altitudes) are considered, the development of the required model is substantially harder. The modeling of each known error type, such as VOR scallops, requires the introduction of additional elements in the system state vector and each of the matrices. As the number of elements increases, it become substantially more difficult to solve the matrix equations.

Making the development of the mathematical model even more difficult is that many errors are not smooth functions, but discontinuities such as the loss of signals from a station due either to the failure of the station or the receiver, or to the inability to continue receiving the station's signal. Switching ground stations can produce transients in the model that may require a long settling time. Furthermore, it is difficult to determine when a station switch should occur based on the Kalman filter's data. While these problems can be handled by special monitoring subroutines in the Kalman filter program that alter the various matrices when a discontinuity is observed, this approach substantially complicates the Kalman filter and may lead to improper operation.

Even if a proper mathematical model could be developed, Kalman filtering exhibits difficulties. If an incorrect sensor measurement were to occur, it would be combined, based on its previous weighting, with the other inputs to form a new state vector and determine new sensor weightings, affecting the system even when correct measurements are again obtained. A software error in the Kalman filter program, undiscovered because it occurs only on rare occasions, can result in improper matrix values which magnify and prolong its effect.

Finally, the numeric computations required for Kalman filtering cause further problems. It may be impossible, given the measured data, to perform the matrix inversions required to solve for the next state; surprisingly, measurements known to be very accurate or with highly correlated errors compound this problem, although they should, in fact, make position determination easier. Even if ill-conditioning could be avoided, the avoidance of round-off errors requires the use of high-precision floating-point arithmetic. This, combined with the computation requirements of the matrix methods used in Kalman filtering (which expand with the cube of the number of elements in the state vector, which in turn determines how well-known error classes can be modeled), makes implementation difficult on processors within the cost, size, and power constraints of many aircraft.

An Alternative Approach

Rather than attempt to modify the Kalman filtering technique to minimize the problems previously discussed, an alternative that more closely matches the characteristics of low-cost microprocessors appears attractive. It is based primarily on logical decisions with limited precision mathematics used in the most frequently executed portions of the program. Rather than use a general filter that attempts to create a model that describes all eventualities, special subroutines that directly address known major difficulties (such as VOR scallops or Omega SID's) are used. This use of "defensive" subroutines, similar to techniques common in Omega navigation systems for handling diurnal variations, takes full advantage of the logical control structure of a digital processor.

Because of its complexity, it is virtually impossible to write a single program that properly anticipates all combinations of sensor errors, and even if such a program could be written and properly debugged, it would be difficult to add or change navigation subsystems when better sources of navigation information, such as the GPS, become available. This problem can be substantially reduced by partitioning the system into a number of parts. Figure 1 illustrates the most logical partitioning with a *tracker* assigned to each navigation sensor (in this example, VOR and Omega receivers and Dead Reckoning), a *combiner* to determine the best estimate of position based on the information from the trackers, and a *navigator* to determine and display deviation from a user-selected course. Because of its simplicity, and since it is identical to the navigation algorithms used in other approaches, the navigator will not be discussed.

To aid in the partitioning of the program, and to allow easy changing of the input navigation systems, it is important that the messages between the trackers and the combiner be as flexible as possible; to minimize program complexity, it should also be simple. The format selected consists of an octuple $(X, Y, S, B, T, K, BLAT, BLON)$ called an *envelope*.

The first five parameters are illustrated in Fig. 2; X and Y indicate the tracker's best estimate of position, the conventional output of a navigation system. S, B, and T define a rectangle that, with high probability, contains the actual aircraft position. Intuitively, it corresponds to the 2-sigma ellipse that results when the normal distributions of the errors in a two-dimensional system are considered. When the tracker determines that unremovable errors are present, it presents a very large rectangle to indicate the lack of confidence in the position estimate.

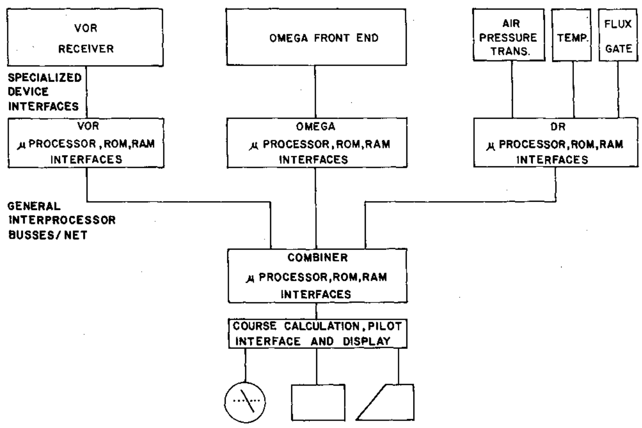

Fig. 1 Block diagram of major system components.

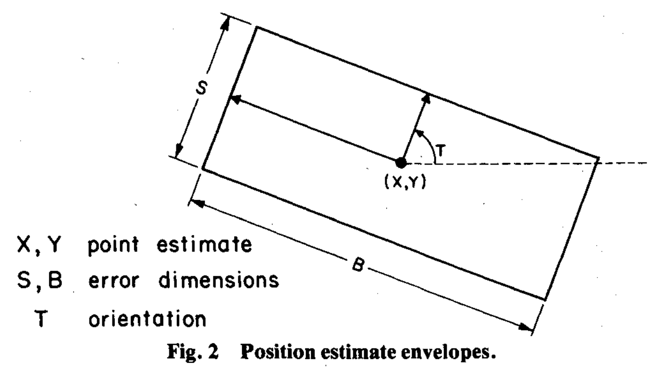

Fig. 2 Position estimate envelopes.

The K parameter is a state flag that defines special attributes of the envelope for use in the combiner, and indicates how the selection of the final position was made when the result envelope is sent to the navigator and the trackers. If this were not done, a tracker may attempt to improve its performance based on a combined position that was based heavily on its last envelope; unless carefully considered, this feedback could make the operation of the tracker, and the total system, unstable due to magnification of incorrect position estimates. Finally, *BLAT* and *BLON* are the locations of the origin of the flat-earth plane used to simplify the computations.

Subsystem Selection

While any source of navigation information can be used as an input to a tracker, best performance of the system depends on their careful selection. Obviously, the more accurate the sensor, the better the operation of the total system. However, techniques with only marginal accuracy, such as VOR theta-theta navigation, can be used provided that the types and effects of errors are understood. In this regard, errors can be classified into two classes. The first class of errors are those which, if an accurate estimate of position were known, would be easy to determined due to an identifiable pattern in the error between the computed and actual positions. This pattern is called the *signature* of the error. For example, in the presence of multipath interference, or scallops, the VOR-derived position for a straight course forms a sinusoidal track (this will be discussed in detail later). The other class of errors are those which have no recognizable signature and are generally regarded as random noise.

Since the primary task of each tracker is not only to use simple filtering techniques to reduce the effects of random errors but also to determine from their signatures what recognizable errors are present and to take corrective action, the navigation sources should be selected to aid in the identification of the error signatures. This can be done by selecting the systems so that, for each type of error, only one out of three or more systems is affected. (When subsystems do not share common effects from potential error sources, their errors are said to be *orthogonal.*)

For example, if the VOR tracker observes a sinusoidal flight path, it is unable to determine if it is the actual track or the signature of multipath scallops. By checking the perceived flight path determined by the Omega and dead reckoning trackers, whose errors for multipath interference are orthogonal to VOR, it is possible to determine whether the scallop signature is present along with its magnitude, permitting its removal. Even though the other trackers are also being influenced by errors, it is unlikely that their effect will produce a signature that matches the VOR scallops'.

Because of their low cost, good availability, and known error signatures, it was decided to use VOR theta-theta navigation, Omega, and a very simple dead reckoning (DR) system consisting of a fluxgate and pitot-static pressure transducers (with about 10% error) in the initial feasibility study. These three sources have good orthogonality in their errors,[8] are not of high quality and therefore provide a "high stress" test of the approach's capabilities, and permit the construction of a very inexpensive system for general aviation.

The Combiner

The heart of the approach is the combiner, which takes the estimate of position envelopes from the trackers and produces a best estimate of position which is then fed to the navigator and back to the trackers. In many respects, it is a direct replacement for a Kalman filter, although it requires its inputs in a different form (an envelope rather than simply a position estimate) and does not attempt to determine or remove sensor errors (which is done by the trackers).

Rather than combine all envelopes in a single operation, the combiner operates on them pairwise. Two different strategies are employed in comparing the envelope pairs—one if they are from relatively independent sources (such as VOR and Omega), and the other if one is highly dependent on the system position estimates (as in the case of DR, which uses the final position estimates to compute a wind estimate[9]). This is because, while the dependent system is very useful as a confidence check in the next position estimate and can be used to identify error signatures, its accuracy can never be better than that of the previous position fixes, which were based on the other sensors.

In the strategy used from independent sources, when the envelope pairs are compared, there are three different cases, as shown in Fig. 3. The first case is when the point position estimate of each envelope lies within the error rectangle of the other. This means that the two envelopes from the trackers are in substantial agreement as to the aircraft's position, and is the expected situation. In this case, the position estimate from the envelope with the smallest S dimension (the smaller side of the rectangle) is chosen, since its tracker indicates that it has the better information or the lesser error. In the case of the top example of case 1, the horizontal rectangle is chosen over the square; its length (B dimension) is reduced by the combiner to match the projection of the other envelope to reflect the better knowledge of the aircraft's position (in this case yielding the smaller enclosed rectangle).

In case 2, one rectangle excludes the other envelope's point position estimate. Again, there is no dramatic disagreement between the trackers, since one of the point position estimates is contained within the error rectangles (and therefore, the best guess of position) of both systems. The envelope with the point position within both rectangles is therefore selected, since it is consistent with both systems. As in case 1, the selected envelope can be reduced by the projection of the other envelope.

Case 3 has the point position estimate of each excluded from the other's error rectangle, and can exist in a variety of forms, the most common of which are shown in the figure. The combiner attempts to apply a number of heuristics to determine a composite envelope. If the envelope pair consists of two elongated rectangles that intersect, as would be the case of a degraded mode position estimate where only a single VOR radial or Omega line-of-position is available, the best composite would have the position estimate at the center of their intersection (as shown by the top example of case 3). Again, in the middle example, the compromise position estimate is located in the intersection of the two rectangles.

The bottom example of case 3 provides a dilemma for the heuristics, since the envelope from each tracker indicates that it is confident of its position estimate, but at least (but probably only) one is in error and has either underestimated

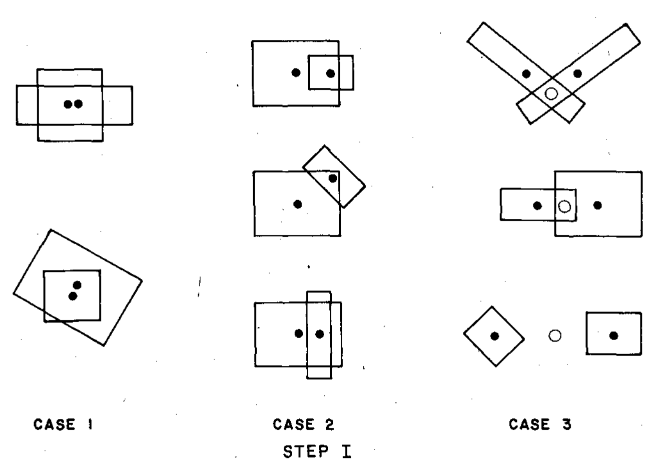

Fig. 3 Different pairwise cases in the combiner.

its accuracy or has experienced some unexpected failure (undetected sensor difficulties or a hardware or software error that caused a computational mistake). The combiner takes two actions. First, it creates a compromise envelope with a point position located between the two point positions, with the exact location based on the ratio of the areas of the two error estimates (placing it closer to the smaller error estimate), much like a standard weighted combination such as in a Kalman filter; and second, it marks the compromise envelope to indicate that a possible difficulty exists. If during other pairwise envelope combining using this compromise another inconsistency is found, the combiner discards the composite and selects one or the other original envelopes if they are consistent in the later combination, and informs the unselected tracker that it has a possible error. This can cause a forced reset of the tracker, hopefully recovering from the problem that caused the original mistaken envelope.

The combining of a dependent envelope also has three cases, similar to independent envelopes, but the actions taken are somewhat different. In cases 1 and 2, if the independent tracker's point position is within the envelope of the dependent tracker, it is chosen as the result, since it has more recent and accurate information. If, in case 2, only the dependent tracker's point position is within the other's envelope, that point position is selected; this would occur when there is a sudden period of poor quality sensor information, and the dependent tracker is needed for "coasting through" the difficulty.

If a complete inconsistency (case 3) is encountered, the combiner checks the flag to determine if this is a compromise of a previous inconsistency. If it is, it attempts to use one or the other of the original envelopes to find a consistent combination. If this fails, the combiner indicates a serious difficulty and produces no position estimate, since all navigation sources are inconsistent.

Before being used by the navigator, one additional bit of processing is required. Since sequential position estimates from the combiner are highly correlated (smooth) only as long as one tracker is being consistently selected, and may have jumps as switches between trackers occur, a smoothing process such as a low-pass digital filter is applied. While this introduces minor lags in switching to a better position estimate, it occurs only infrequently and helps reject spurious mode changes.

Implementation

The combiner can be implemented with little difficulty on available microprocessors. For enroute operation, it need only be scheduled every 10 s, corresponding to the update cycle of Omega; for terminal and approach operations, or when sources with higher update rates are used, it can be scheduled more frequently without overburdening the microprocessor. If case 3 operations, which occur only when difficulties in the trackers exist, are ignored, the worst case path through the combiner requires about 500 different operations.[9] Over half of these are assignments or fixed-point additions and subtractions, and only about 30 use trigonometric functions; no matrix operations are required. Case 3 approximately doubles the worst case instruction count.

In most instances, high precision or floating-point operations are not necessary. A simple reorigining of the lat/long position values (consisting of subtractions) can reduce the required accuracies to under 16 bits. Trigonometric functions need only have 8-12 bit accuracy, and can be obtained using either table lookup (possibly with linear interpolation) or by the CORDIC algorithm.[10]

Trackers

While the operation of the combiner is relatively independent of the source of navigation information, each tracker in the system is designed specifically for a given source. They are responsible for the management of the source's sensors (including selection of necessary stations) and removal of characteristic errors based on their signatures. For the initial feasibility study, trackers for VOR theta-theta and Omega radionavigation and DR were developed.

VOR Tracker

The input to the VOR tracker is digitally decoded phase difference information from one, two, or three different ground stations. This may be supplied by three different receivers or by a single receiver that is time-division multiplexed among the different stations; in any case, station selection is determined by the tracker and requires no pilot intervention.

Three different types of errors predominate in the VOR systems: bias errors, multipath interference, and random noise. Bias errors are caused by misalignment of the ground station or unanticipated phase shift differences in the receiver and digital decoder; the former is minimized by constant monitoring and control of the VOR station's radiation pattern. Receiver biases appear as the equal rotation of all ground station radials, and a simple calculation based on the station geometry can determine the amount of error, although this was not implemented in the initial study.

The more common VOR error is caused by the uncontrolled reflection of the ground station's radiation by objects near the station. Although careful site preparation reduces this problem, it can produce errors in the location of radials of as much as 6 deg. The reflected signal creates an interference pattern with the correct signal, causing a distortion in the amplitude of the signal that indicates the aircraft position in the phase comparisons. The amount of distortion is a function of the VOR frequency, the relative positions of the VOR, aircraft, and reradiator, and the reradiator's percentage of reflection. The general effect is to distort the radial so that it forms a sinusoidal pattern about its true position, generally referred to as a scallop or bend.[11] (This is what causes the "windshield wiper" effect in a lightly damped VOR indicator as a course is being flown or an intersection approached.)

If the flight path is not near the VOR (which can easily be determined by the tracker, causing a different set of routines to be used) and is generally following a straight course, then for moderate periods of time the true azimuth of the aircraft from the station should change linearly. However, if the radials are scalloped, the received azimuth will be sinusoidal about the true azimuth, varying with a period of 0.01-10 Hz (see Fig. 4). To remove the scallop error, it is necessary to determine its amplitude and period; then, a simple prediction of the error and subtraction from the received azimuth can be employed.

This can be done by using a least-squares estimator to determine the true azimuth from the distorted received radials. It requires the sampling rate to be a function of the sinusoidal distortion, with about 16 samples used, each at about one-eighth of the period of the distortion. Although this method requires more samples than theoretically necessary, it also removes many of the random noise errors by acting as a high-pass filter. If the intermediate terms used by the estimator are stored, the next estimate can be produced by subtracting the oldest intermediate term, adding the intermediate term produced from the latest sample, and performing a handful of calculations, all well within the capabilities of a microprocessor.

In steady-state straight flight, the least-squares estimator provides refined azimuth values without the lag characteristic of other low-pass filters. However, if there is a change from steady-state, the estimates will lag until the past intermediate terms have been substantially removed from the estimator's memory. While this will occur naturally with time (20 s for the lowest scallop frequency, 0.01 Hz), it is possible for the tracker to force the rapid dumping of intermediate terms if a

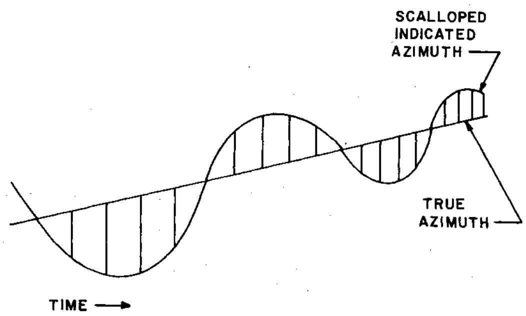

Fig. 4 Least-squares tracking of a VOR signal.

drastic change in course is determined. Both heuristics looking for these changes and information from the other navigation sources and the combiner can be used to determine if dumping should occur and how much should be dumped.

Construction of an envelope requires the reception of two VOR stations, whose azimuth intersection is the point position estimate. The error envelope is constructed from deviations about the estimated radial based on the magnitude of the scallop distortion; a minimum error of 1 deg is used to insure that the envelope dimensions are conservative. Periodically, the tracker also forms envelopes between each of the current VOR stations and a VOR station from the third receiver (initially selected because of promising geometry). If any of these envelopes are substantially better, their VOR stations become the new current pair, providing automatic station selection based not only on geometry but also signal quality.

If only one station is being received, or if the stations and the aircraft are approximately colinear, a special envelope is formed. This envelope has a B parameter (longer side) arbitrarily long, its S parameter (shorter side) based on the amplitude of the single VOR errors, the point position on the estimated radial at a point based on past information, and a special flag set in the envelope. This allows the combiner to use the degraded information to improve on the estimates from the other systems.

Omega Tracker

The Omega tracker is similar to the VOR tracker, except that the defensive code acts against errors caused by changes in the transmission waveguide formed by the Earth and the ionosphere from transitions between daylight and night (diurnal variation), or SID's caused by sunspots. Both have definite characteristic signatures, diurnal variations producing a pattern based on the aircraft's position relative to the Omega ground stations and the zone of twilight, and SID's only distorting the portion of Omega signal's path in sunlight.[12] While the errors of VOR are associated with the location of the aircraft relative to the ground station and reradiator, the errors in Omega are basically constant over large areas but vary with time.

While it is possible to compensate for many Omega errors by complex calculations called *skywave correction factors,* there are alternative ways to determine the current Omega error, which will be valid for a large area and a moderate length of time, and apply it to the received information. These can then be periodically altered based on known trends until another accurate estimate can be made. At one time, it was proposed that special Omega monitoring stations be deployed, each determining the skywave corrections for an area and forwarding it to the using aircraft[13]; this plan has been abandoned because of the resulting cost and complexity.

However, if an accurate estimate of position is available independent of Omega, it is possible to determine the error and form a *pseudo-skywave correction* (PSWC). The accuracy of the PSWC is dependent on both the uncorrelated noise in the Omega signal and the error in the position estimate (generally from the VOR tracker), and should only be calculated when an accurate position estimate is available.

Since extensive self-contained techniques for recognizing and minimizing propagation disturbances are not required in the Omega tracker because of the use of the PSWC technique, it is substantially less complex (and therefore less expensive) than any available Omega receiver. All that is required is a very low-frequency amplifier and a digital decoder, based on phase-locked loop techniques to provide good recovery in the presence of high noise.[14]

In addition to the tracking of the received Omega signals and the updating of the PSWC, infrequently, the Omega tracker must determine the necessary conversions between the local flat-earth and the more accurate round-earth models. Although this requires a substantial amount of computation, it need be performed only every few minutes.

Dead Reckoner

The dead reckoning tracker is extremely primitive when compared to standard air data systems, being implemented using medium accuracy (within 10 knots and 10 deg) sensors. It determines its wind estimates from the difference between measured velocity and the change in the best estimate of position, making it dependent on the other trackers (and requiring a different combiner strategy). Special actions must be taken when a dramatic turn (e.g., over 45 deg) is detected to avoid having the effect of sensor errors doubled because of their inclusion in the estimated wind.

Simulations

Because of the high degree of interaction between navigation sources, especially in the combiner, and the ability for a tracker to drastically alter its processing methods as input conditions require, it is virtually impossible to mathematically analyze the operation of the system and determine its probable error magnitudes. However, the algorithms for the combiner and trackers can be interfaced to a special simulation system that schedules the execution of the combiner and trackers, provides inputs such as would be available for the various sensors, and monitors and analyzes the system's performance. Such a simulator has been written in FORTRAN and runs on a Cyber-175 computer.

These simulations were used to debug the tracker and combiner algorithms and determine the effect of changes to the technique, in addition to determining the feasibility of the approach. In some cases, complete versions of the routines were not used; the VOR tracker does not deal with bias errors and the Omega tracker does not recognize SID's. The models used in the simulator do not produce these effects, although both the models and the trackers can be easily modified to handle them. By its nature, the complete combiner was necessary.

Since much of the system is concerned with the removal of types of errors, it is important that the simulation contain realistic models for the navigation systems. In fact, it is best that the models contain errors whose magnitude is higher than normally expected so that the system is properly exercised. (If only small, or no, errors were considered, any navigation system would exhibit satisfactory performance.) The model includes two VOR stations located about 30 n.mi. (54 km) apart and 10 n.mi. (17 km) north of the flight path. Reradiators near the VOR's produce scallops with up to 5.7 deg of error with random noise of 0.15 deg rms superimposed. (Generally it is assumed that 95% of the time actual VOR errors will be less than 1.9 deg.[1]) A random noise of 2.5 centicycles rms, corresponding to 0.325 n.mi. (0.6 km), is used for Omega; this is about three times the error claimed for corrected Omega.[11] Finally, a 10 knot (19 km/h) and 10-deg error is added to the DR sensors.

Figure 5a shows the flight path used for the simulations, and the estimates of position provided by the system, with no

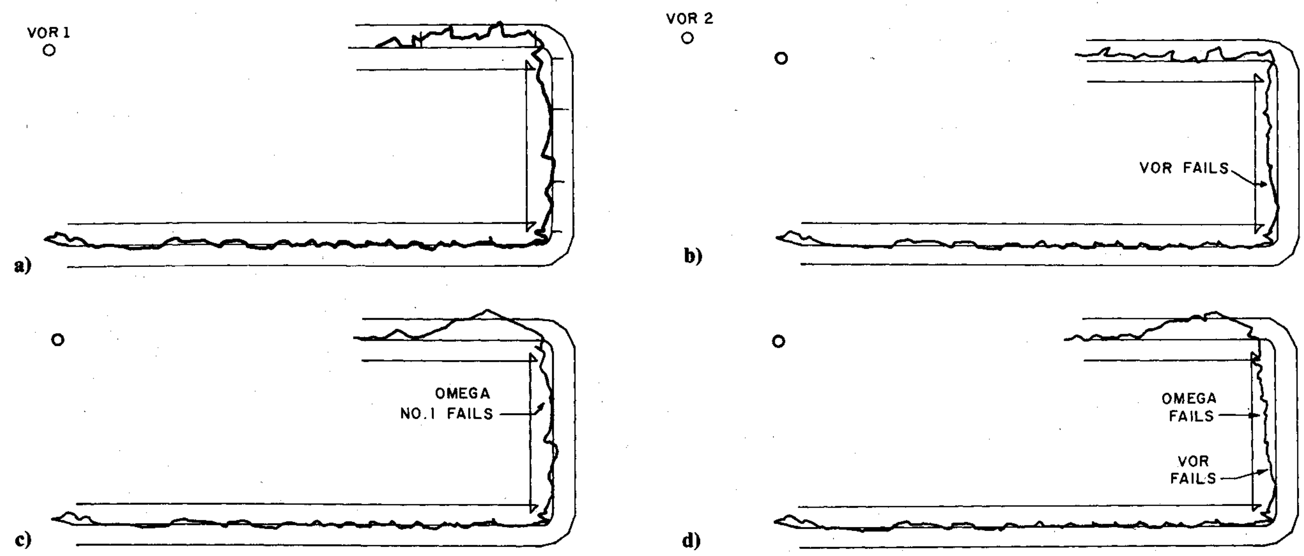

Fig. 5 Simulation position estimates and effects of sensor failures: a) normal, b) VOR failure, c) Omega failure, d) Omega and VOR failure.

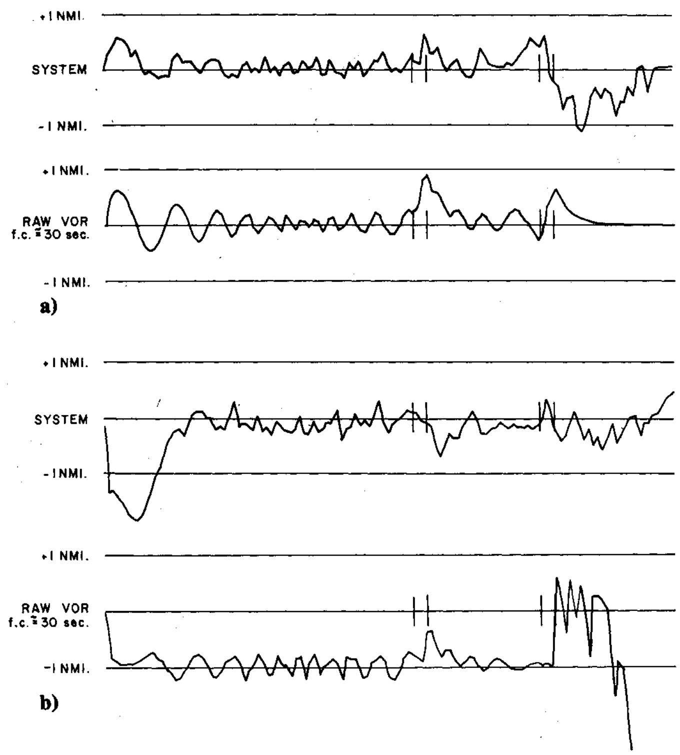

Fig. 6 Strip chart portrayal of a) crosstrack and b) alongtrack errors.

smoothing or filtering applied. The scale on either side of the flight path is 1.0 n.mi. The flight starts at the left, south of VOR 1, at 2400 GMT when the Hawaii-Trinidad Omega phase, necessary for navigation at the assumed central Illinois simulation location, is making a rapid transition. The aircraft accelerates instantly from zero to about 100 knots, producing a transient in the behavior of the system. Although special subroutines could be used to better handle this startup, it does produce a good test of the transient response of the system.

While Fig. 5a indicates that the absolute crosstrack error generally remains less than 1 n.mi. (and often less than 0.3 n.mi.), it is impossible to determine the alongtrack error. Figure 6 shows the same simulation results in the more familiar "strip chart" form with the location of the two turns shown by the hash marks on the zero line. For comparison, the error in position for VOR navigation using only a standard filter with a constant of 30 s is shown. For crosstrack errors, after the initial startup, the system's estimate of position is comparable to filtered VOR. After the second turn, the filtered VOR error goes to zero since the flight path is colinear with the two VOR's, while the only system crosstrack error over 1.0 n.mi. occurs. This happens because the code that rapidly dumps past observations in the VOR least-squares estimator has been disabled to see its effect; with it enabled, the error is substantially less.

While there is little improvement in crosstrack performance, there is substantial improvement in alongtrack error. After the initial startup error (which could be removed with special routines), the system settles to an error generally less than 0.4 n.mi. Because of the large filter constant necessary for acceptable crosstrack error on the raw VOR navigator, a lag in position of about 1.0 n.mi. results. Furthermore, while the system shows good performance on the colinear segment after the second turn, the VOR navigator has an unbounded alongtrack error.

Degraded Operation

Figure 5b illustrates the effect of the loss of a VOR station (in this case VOR 2, the one on the right); Fig. 5c, the loss of the Norway Omega station (allowing only one line-of-position to be determined); and Fig. 5d, the loss of both in rapid succession. The loss of a VOR station would generally be caused by an obstruction (including the horizon) between the aircraft and the station, while Omega stations are subject to a variety of unanticipated outages. In all three cases, the system performs reasonably; in fact, it appears to perform better after the failure of VOR 2 than in the normal case! This is because the loss of VOR 2 forces the tracker to execute the otherwise disabled code to dump intermediate results during and after a turn. With that code enabled, the results are comparable. The overshoot after the second turn following Omega failure can also be reduced by forced dumping. With conventional navigators, and even with a Kalman filter combiner, the determination of position in the case of both VOR and Omega failure would be impossible.

Conclusions

Although not extensively tested, the approach described provides at least a mininal tolerance to a wide variety of errors—sensor failures, failures in the trackers' microprocessors, and undetected program bugs in the trackers—without special programs or hardware. Alteration to the combiner's structure, without major changes to its technique, can improve its tolerance to unexpected failure of its software. Furthermore, only first-order errors, such as VOR scallops or Omega SID's, were modeled; the assumptions used for the dead reckoner, while representative of those expected from available sensors (in fact, errors somewhat greater than from available low-cost sensors), do not specifically match any particular sensor system. Incrementally better performance can be expected as more types of errors, including those which occur infrequently or have only minor effects, are considered by the tracker programs.

While the use of low-cost sensors, such as VOR and a simple Omega receiver, airspeed, and heading transducers, was made in the initial feasibility study, this was only to provide a better test of the approach under stress. If more accurate sensor information, such as DME or the GPS, were available to the system, better estimates of the errors in the other sensors and more accurate position estimates could be made, even if the more accurate sensor information was subsequently lost, since the accuracy of the other sensor's error models benefits from the use of more accurate information.

A number of capabilities that require special processing with other approaches occurs with no additional effort. For example, since the final best estimate of position is generally that of the most accurate system, if a microwave landing system (MLS) receiver were included as an additional subsystem, as the terminal area is approached and the MLS error envelope becomes smaller than that produced by the other navigation sources, automatic transition from enroute RNAV to MLS terminal area navigation results.

It can be easily seen that the logical method of combining navigation information sources can produce acceptable position estimates even in the presence of large sensor errors without requiring extensive onboard computer systems. Unlike matrix techniques such as Kalman filtering, its instruction and data resolution requirements provide a good match to available low-cost microprocessors. The use of trackers that communicate using a simple, fixed protocol with the central combiner allows the replacement or addition of navigation sources without substantial modification. Since the trackers are independent programs, possibly running on different processors, their programs can be structured as required by the sensor and its error model rather than being forced into a common mathematical mold, simplifying their development while increasing their flexibility.

Acknowledgments

This work was supported in part by the National Aeronautics and Space Administration's Langley Research Center under Contract NAS1-15145. The computer time for the simulations was provided by the Research Board of the University of Illinois at Urbana-Champaign.

References

[1] Tyler, J.S., Brandewie, D.M., Heine, W., and Adams, R., "Area Navigation Systems: Present Performance and Future Requirements," *Navigation: Journal of the Institute of Navigation,* Vol. 22, Summer 1975, pp. 95-111.

[2] "Approval of Area Navigation Systems for Use in the U.S. National Airspace System," Advisor Circular 90-45A, Department of Transportation, Federal Aviation Administration, Feb. 1975.

[3] Vass, E.R., "OMEGA Navigation System: Present Status and Plans 1977-1980," *Navigation: Journal of the Institute of Navigation,* Vol. 25, Spring 1978, pp. 40-48.

[4] Kalman, R.E. and Bucy, R.S., "New Results in Linear Filtering and Prediction Theory," *Transactions of the ASME, Journal of Basic Engineering,* Ser. D., Vol. 83, March 1961.

[5] Kayton, M., "The Navigation Equations," *Avionic Navigation Systems,* edited by M. Kayton and W.R. Fried, Wiley, New York, 1969, pp. 13-54.

[6] Farrell, J., *Integrated Aircraft Navigation,* Academic Press, New York, 1976.

[7] Bryson, A.E. and Bobick, J.C., "Improved Navigation by Combining VOR/DME Information and Air Data," *Journal of Aircraft,* Vol. 9, June 1972, pp. 420-426.

[8] Hollaar, L.A. and Cannon, M.J., "Multi-system Position Determination: An Innovative Approach to Low-Cost Area Navigation," *Proceeding of the AIAA/IEEE 2nd Digital Avionics Systems Conference,* Nov. 1977, pp. 230-237.

[9] Cannon, M.J., "A New Approach to Multi-Sensor Navigation System Design," Ph.D. Dissertation, Department of Computer Science and Aviation Research Laboratory, University of Illinois at Urbana-Champaign, Jan. 1979.

[10] Volder, J., "The CORDIC Trigonometric Computing Technique," *IEEE Transactions on Electronic Computers,* Vol. EC-8, Sept. 1959, pp. 330-334.

[11] "VOR/VORTAC Siting Criteria," Handbook 6700.11, Department of Transportation, Federal Aviation Administration, Aug. 1968.

[12] Swanson, E.R. and Levine, P.H., "Automated Omega/VLF Monitoring and Forecasting for Air Traffic Safety Enhancement," *Institute of Navigation National Aerospace Symposium,* April 1978, pp. 71-78.

[13] Hastings, C.E. and Comstock, A.L., "The Optimum Implementation of a Differential Omega System," *Navigation: Journal of the Institute of Navigation,* Vol. 22, Spring 1975, pp. 76-80.

[14] Cox, D.B. Jr., Harrington, E.V. Jr., Lee, W.H., and Stonestreet, W.M., "Digital Phase Processing for Low-Cost Omega Receivers," *Navigation: Journal of the Institute of Navigation,* Vol. 22, Fall 1975, pp. 221-234.