

ON ESTIMATING MAP MODEL ERRORS AND GPS POSITION ERRORS: (APPLYING MORE SCIENCE TO THE ART OF NAVIGATION)

by Peter KIELLAND¹ and Tony TUBMAN²

Paper presented at the Institute of Navigation's annual Technical Meeting, San Diego CA, 26 January, 1994.

Abstract

In order to decide whether a desired manoeuvre can or cannot be safely undertaken, a prudent navigator must be aware of *both* the current spatial uncertainty of his vehicle's positioning system and the spatial uncertainty of the navigational map model being used to depict the theatre of operations. From this safety to navigation perspective, knowledge of data accuracy is as important as the data itself. This paper discusses the Electronic Chart (EC) implications of both GPS vehicle positioning errors and the relatively large data modeling errors specific to bathymetric map models (charts). It proposes and demonstrates software solutions which statistically evaluate both of these spatial uncertainties and graphically integrates the two stochastic models within an EC environment. The paper also documents an experiment carried out by the Canadian Hydrographic Service, designed to insure that real-time DGPS users compute statistically valid position error estimates. The experiment ground truthed the position *error estimates* obtained using a conventional real-time error analysis of pseudo-range redundancy. Using this ground-truth information, an improved pseudo-range error model was empirically determined. The new pseudorange error model is continually updated using the estimated pseudo-range variances computed by the Novatel GPS receiver rather than applying the constant a priori pseudo-range variance typical in least-squares adjustments. This dynamic range error model effectively reduced the statistical bias between the observed errors and their predicted error estimates. The improved range error model also significantly improved the performance of the position solution. All DGPS positions computed by the modified software had a positional accuracy of better than 0.5 metres.

¹ Canadian Hydrographic Service, Ottawa, Ontario, Canada.

² Nortech Surveys Canada, Calgary, Alberta, Canada.

INTRODUCTION

Safe navigation is both a science and an art. The *science* of navigation is exemplified by the measurements and computations which lead to GPS position estimates. The *art* of navigation consists of *using* these position estimates to safely and efficiently accomplish a productive navigational task. Navigation occurs when a vehicle's observed present position and course history are used to plan future motion relative to a map model that depicts the real world. Safe yet efficient navigation occurs when a navigator evaluates and balances the *risk* involved in any planned course. Sustaining a high level of navigational risk will eventually result in an accident, thus curtailing productive work. At the other extreme, truly minimizing navigational risk implies a timidity which also limits profitable activity. This paper documents new techniques being developed to statistically evaluate navigational risk. It also illustrates how these computed indicators of spatial risk can be represented graphically and presented to navigators in an easily perceived format. By quantifying navigational risk more *statistically*, a navigator's intuitive ability to correctly evaluate spatial risk can be greatly augmented.

The algorithms discussed in this paper are intended to operate within an "Electronic Chart" environment. Whether used for land, sea or airborne applications, an Electronic Chart (EC) is simply a real-time computer program which graphically integrates digital map information with observed vehicle positions. In its most basic nautical form, the EC automatically plots a vessel's position onto a digital image of a map model. The first generation of Electronic Charts can only display bathymetric models which are simply images of traditional paper charts converted into a computer display format. The real power of the EC will be realized when data analysis tools permit navigators to "see" further into our growing digital database of spatial information than is possible using traditional cartographic renditions of the data. This paper describes new analytical tools which extend our vision into that mass of numbers and extract the risk information needed to practice safe navigation.

THE TWO PARTS OF THE PUZZLE

A navigation *system* is composed of both a map model and a vehicle positioning system. To quantify the risks involved in using this system, two stochastic elements must therefore be modeled:

1. The spatial uncertainty of the map model being used for planning the route.
2. The spatial uncertainty of the vehicle's current position used to follow the planned route.

In the GPS navigation literature to date, risk analysis investigations have been largely concerned with the first element: the positional uncertainty of the vehicle that is being navigated. This groundwork was initiated primarily for airborne

applications and has evolved under the general heading of "GPS Integrity Monitoring". The reason that airborne navigation applications can legitimately ignore map model errors is that aircraft (hopefully) have very limited interaction with terrestrial map models. Runways are neatly bounded bits of flat, man-made terrain which can be easily and accurately modeled within the GPS frame of reference. Following safe glide paths onto runways thus becomes purely a matter of providing sufficiently accurate and reliable vehicle positioning to guide the aircraft onto these precisely located targets. However, once we leave the world of airport runways and start navigating with respect to the seafloor the error modeling situation changes drastically.

QUANTIFYING THE UNCERTAINTY OF BATHYMETRIC MODELS

Seafloor terrain is inherently difficult to survey. This difficulty stems from the high cost of measuring sufficiently dense depth soundings to fully define its shape. Economics dictate that the seafloor which lies between discrete hydrographic sounding profiles must be *interpolated* from the sparse measured data. In areas of rough geomorphology, this can result in very large depth interpolation errors in the map models ultimately used by navigators. An example of this is the recent grounding by the QUEEN ELIZABETH II on a shoal that had gone unsampled between the hydrographic survey profiles. Dangerous interpolation errors were obviously present in the charted depth contours drawn through the survey data yet this depth uncertainty was not apparent to the Captain.

Continuous depth contours derived from discrete soundings are contaminated with a *composite* of interpolation errors and instrumental errors. Instrumental errors are inherent to the survey equipment used to measure soundings. Interpolation errors are inherent to the density of the soundings (the survey line spacing) with respect to the complexity of the seafloor. To fully quantify the uncertainty of a surveyed bathymetric model, both instrumental and interpolation errors must be continuously modeled for all locations within the surveyed area. Only then can the hydrographic data be safely exploited in an Electronic Chart environment.

Navigators run the risk of getting into trouble when an EC integrates images of traditional paper charts with the highly accurate vessel positioning provided by differential DGPS. The accuracy of GPS positions can be well predicted and displayed in real-time (the objective of the experiment described later in this paper). However, the accuracy of the bathymetric model portrayed by a traditional paper chart is really unknown and unstated on the document. If this image is digitized and viewed on the screen of an Electronic Chart, there is strong temptation for navigators to conclude that the bathymetric model being viewed is as modern as the rest of the EC package. Depending on a navigator's background, he might easily conclude that the chart image portrays a totally reliable bathymetric model free of instrumental and interpolation errors. More cautious (realistic) navigators will treat all bathymetric data as being "only approximate", however, this vague feeling about map model fidelity does little to help in actually choosing a course that is safe yet economically efficient.

Since the spatial real uncertainty of bathymetric models varies tremendously from place to place, Hydrographic Offices must do their best to objectively quantify its uncertainty at *all* locations and provide that information to users of the data. Unfortunately, the vast majority of our planet's seafloor is only surveyed with widely spaced soundings. These measurements are also sometimes quite inaccurate. The economic realities of carrying out new acoustic swath surveys dictate that this problem will remain with us for the foreseeable future. Since we are stuck with these often very approximate bathymetric models, their spatial uncertainty must be statistically modeled and made available to navigators.

The hydrographic community is currently addressing this very real problem. A standard algorithm for statistically describing the spatial errors inherent to bathymetric data is being developed and implemented by members of the International Hydrographic Organization (IHO). These spatial statistics are computed using a geostatistical depth interpolation algorithm which *also* predicts the depth estimation errors inherent to each point on the interpolated bathymetric model. The software being developed employs a classical grid interpolation algorithm called "kriging" and it is being made available as a Public Domain program called "IHOstat". An interpolation algorithm such as kriging interpolates a matrix of depths which are then typically used to generate graphical interpretations of the survey data such as depth contours or 3D surface models. The mathematics of kriging are somewhat complex and outside the intent of this paper. A brief overview is however appropriate here.

The kriging algorithm in IHOstat is essentially a least-squares depth interpolator which uses the variogram as a weighting function. The variogram is a useful geostatistical indicator of bottom roughness which is computed from a population of surveyed depths. A variogram graphs the variability of all possible depth differences in the data sample against the horizontal distance separating each pair of soundings used to compute the differences. This relationship provides a statistical basis for predicting the uncertainty of an interpolated depth at any given distance from a measured sounding.

Variograms are computed within small local regions throughout the data set. This creates a "map" of changing bottom roughness which then is used to interpolate of the overall bathymetric model. At each desired grid location, kriging uses the local variogram model to assign an appropriate weight to each sounding in the neighbourhood of the grid point as it is used to interpolate that depth. The farther away a sounding is from the grid node, the less its statistical weight will be in the least-squares interpolation. Conversely, if the grid node being interpolated is exactly coincident with a measured depth then that sounding's weight will be 100% and the grid node takes on its exact value (i.e. kriging "honors" the data). By fully exploiting variograms, kriging permits the spatial variability of the seafloor to be used to control the gridding process. The stochastic nature of kriging permits the IHOstat algorithm to interpolate a grid of depths from the survey data and *also* to estimate a standard deviation for each of those depth estimates.

Figure 1 illustrates these two IHOstat products. The lower surface is a "bathymetric surface" of gridded depths interpolated from hydrographic survey data. The upper surface is its corresponding "stochastic surface" of depth error estimates. The height of each grid node on the stochastic surface represents the estimated




FIG. 1.- Different views of the bathymetric and stochastic surfaces.

standard deviation of the corresponding depth estimate on the bathymetric surface. For graphical clarity, the upper surface has been exaggerated 5 times (i.e. it is a 5 sigma stochastic surface). Each of the estimated depth interpolation errors on this surface is a function of the horizontal distance from the grid node to the nearby measured depths that were used in the interpolation: the closer to a sounding a grid node is, the lower its estimated uncertainty.

The contour map in Figure 1b illustrates this more clearly using a different stochastic surface. We can see "ridges" of depth uncertainty which grow (governed by the local variogram models) and peak midway between the measured depth profiles. The maximum predicted interpolation error within any local area is thus a direct function of the sounding line spacing chosen for the hydrographic survey. Since typical hydrographic surveys use a fairly constant line spacing over large areas, the predicted interpolation errors also are sensitive to changes in the roughness of local seafloor morphology: the smoother the seafloor morphology, the lower the expected interpolation errors and vice versa. Since variograms have modeled changes in seafloor roughness, IHOstat's estimates of interpolation error can correctly portray this trend. We can see this correlation between seafloor roughness and estimated interpolation error by comparing Figure 1b (contours of a stochastic surface) with Figure 1c below it (a perspective view of the bathymetric surface it describes).

IHOstat assures the statistical validity of the computed stochastic surface using an automatic process to "calibrate" the predicted interpolation errors. The procedure involves extracting a subset of the surveyed soundings to act as calibration points, then using the remaining soundings to interpolate depths at their exact locations. The differences between the surveyed and interpolated depths (the real interpolation errors) are then compared to the *predicted* interpolation errors for those points. Prior to proceeding with the actual gridding, statistical tests are applied to determine if a bias exists between the real and predicted errors. If a bias exists then the interpolation error models (the local variograms) are adjusted proportionally to the observed bias.

Data "integration" is the task facing a Hydrographic Office when bringing together overlapping data from diverse surveys to create the single bathymetric surface presented to navigators. The instrumental errors of these different soundings are highly variable and depend on the survey technology that was used to collect each data set. How can an optimal bathymetric surface be determined from all this non-homogeneous data? An extension to IHOstat's kriging algorithm has recently been implemented which effectively solves this problem. An estimate of each sounding's position and depth uncertainty is first made using a priori knowledge about the performance of the survey technology that was used to measure it. These instrumental error estimates are then propagated through the kriging equations. By using instrumental error estimates in its kriging process IHOstat takes on two interesting characteristics:

- 1) It produces a fully descriptive stochastic surface (i.e. one that describes *both* the instrumental errors and the interpolation errors).
- 2) Diverse data sets can be rigorously *integrated* into a single bathymetric model. This data integration is possible since each sounding's instrumental uncertainty is taken into account when determining its weight in the least-

squares interpolation of each grid node. Less reliable depth measurements from older surveys are thus appropriately de-weighted when creating the bathymetric model from diverse survey data.

IHOstat evaluates the composite effect of both instrumental and interpolation errors at every location within the surveyed area. We can use this spatial uncertainty information to provide a powerful risk analysis tool to navigators. Real-world navigators are not interested in trying to interpret numerical confidence statistics. The information that navigators need to see on their EC screen is the *implications* of that data uncertainty. It is *horizontal* guidance that is required to steer a safe course. Thus the key to displaying an easily understood portrayal of spatial risk lies in transforming vertical uncertainties in the bathymetric model into horizontal confidence zones which can aid in planning safe routes.

One means of graphically delimiting these confidence zones is to draw a series of depth contours to portray a series of statistically biased bathymetric models. Statistically biasing the bathymetric model produces deeper and/or shallower scenarios for the seafloor model with respect to the most *probable* one. Statistical biasing is accomplished at each grid node by simply adding or subtracting the estimated standard deviation of the grid depth from its most probable depth estimate (the interpolated value itself). A biased bathymetric model is thus computed by adding or subtracting the total stochastic surface grid from its corresponding bathymetric surface grid. Figure 2 illustrates how contouring a series of these biased models might appear as an overlay view within an Electronic Chart environment. Each series of five contour lines was drawn using the most probable bathymetric surface and four statistically biased ones. The center line of this "contour envelope" represents the most probable location of the contour. The two contours on either side of it were drawn by IHOstat using gridded models that were biased by one and two standard deviations respectively towards both deep and shallow interpretations of the data. Together these five contours form a 2-sigma confidence envelope inside of which the true contour will lie with a 95% confidence.

The actual width of the contour's confidence envelope is determined by the magnitude of the stochastic surface. As the model's uncertainty increases, the depth differences between the four statistically biased bathymetric models also increases. This in turn causes the confidence envelope to widen. On the other hand as uncertainty in the model decreases, the five contours which form the envelope will begin to overplot and appear as a single line. The actual width of the envelope displayed on the EC is also sensitive to the changing slope of the seafloor and of course to the viewing scale selected by the navigator.

Figure 2 was computed by IHOstat using soundings surveyed by the Canadian Hydrographic Service at a scale of 1:10 000. If those soundings had been of different accuracy and/or different density, then the width of the contour envelopes in Figure 2 would have also varied accordingly. Since the width of the envelopes is tied directly to the viewing scale selected by the navigator, survey data cannot be abused by an over-confident navigator. Whether the shoal in Figure 2 is viewed at 1:1 000 or 1:1 000 000, its inherent depth uncertainty is correctly scaled onto the EC's display. Since the 2-sigma positioning uncertainty of the vessel's position system has *also* been overlaid onto Figure 2, the image portrays *all* the spatial information needed to take safe calculated risks when passing near the 10 metre shoal it depicts.




FIG. 2.- Integrated graphical presentation of the two stochastic models.

OTHER APPLICATIONS OF GEOSTATISTICAL MODELING

We have seen how geostatistical modeling can produce depth contours which transmit both depth *and* depth uncertainty information to navigators. There are five other practical advantages which can be derived from the technique:

1) Many current EC's are based on raster chart images. These rasters are 'dumb' due to the fact that the value of each pixel in the image represents only the *color* of a scanned paper chart. In contrast to this, an interpolated bathymetric model is a matrix of *depths*. When these interpolated depth values are mapped onto a Color Look Up Table (LUT), the depth matrix is transformed into an "intelligent" raster image. The "intelligence" of this image springs from the fact that each pixel (grid) value represents *both* a colour and a location in 3D space. Traditional chart compilations are perceived as a finite set of symbols, numbers and lines which require a rather cerebral analysis by the viewer. On the other hand, pixel by pixel mapping of depths onto colors is perceived in much the same way as a photographic image. As an optional view in an EC environment, this depth image can effectively transmit very detailed seafloor information to the navigator without causing distraction from navigational tasks.

2) The interpolated matrix of depths can be stored in a far more compact data file than the survey data files from which it is derived. Like any natural image, the bathymetric depth image contains many redundant values clustered together. This permits an interpolated depth image to be compressed using standard compression algorithms developed for general computer graphics applications. The IHOstat grid files can in fact be treated exactly as TIFF image files and processed by standard

image compression/decompression programs. The one constraint here is that the image compression algorithm used should not be "lossy" (i.e. the gridded depth values must be restored exactly when the file is compressed and decompressed). The pixel color/depth resolution must also be sufficient to resolve the full range of surveyed depths encountered in the bathymetric model. This is not a real problem since 32 bit computer imagery can easily provide decimeter resolution of even mid-ocean depths. A bathymetric model's inherent ease of data compression will permit EC navigators to carry tremendous amounts of depth information with them. It would be neither practical nor useful to carry around a massive digital database containing all the surveyed soundings. A bathymetric model can however be easily interpolated to approximately the same density as the soundings in the original surveys. Compression of this matrix data can thus provide an efficient means for navigators to directly view the sounding resolution of the original surveys.

3) In sparsely surveyed areas of the world (i.e. most of it), the computed depth image can easily be interpolated to a much greater density than that of the original surveys. Such dense imagery provides a number of desirable graphical viewing possibilities (discussed below). However, displaying "manufactured" additional depths to navigators would normally be considered dangerous. Navigators viewing this dense depth image (or contours extracted for it) might easily conclude that the measured hydrographic soundings were much denser than they really are. In an EC environment this could easily create dangerous over-confidence in the fidelity of the bathymetric model. This is where the depth image's computed stochastic surface becomes an indispensable aid to safe navigation. It provides the reliability information needed to insure safe usage of the model being viewed. No matter what image density has been interpolated from sparse existing data, the model's computed stochastic surface will correctly estimate both the magnitude and location of its spatial uncertainty. Thus, even for the many areas of the world where only sparse sounding coverage exists, high density depth matrices can be safely interpolated and exploited in an EC environment.

4) Figure 2 illustrates only one method of portraying the bathymetric and stochastic surfaces. Other graphical presentation techniques exist which EC manufacturers could easily implement. Due to its grid format, the information contained in the depth image has extremely useful viewing characteristics. One practical advantage is computational speed. The *gridding* process itself is very long and CPU intensive. For example it took over 20 minutes on a 386 notebook PC to compute the five bathymetric models used to create Figure 2. Approximately 3 500 surveyed depths were used to interpolate the grid to approximately 3 times the sounding density of the original survey. Interpolating large survey data sets will therefore require major computing resources and overnight processing runs. However, all the *intensive* number crunching is done by the Hydrographic Office that collects and integrates the depth data ... not the navigator who *uses* the data.

Once all the survey data has been integrated into a model, extracting arbitrary contours or generating 3D views from the depth matrix is very fast, even with quite limited computing resources. For example, once the kriging was completed, the different graphical representations shown in Figures 1 and 2 can be extracted from the bathymetric model and drawn to the PC's screen in less than 10 seconds. This ability to rapidly change the way data is viewed provides some interesting possibilities for exploiting the Electronic Chart:

- Since arbitrary contour intervals can be quickly extracted from the model, real-time, tidal reductions could be easily applied to the bathymetric model every few minutes and the screen refreshed with new depth contours. Automatic tidal reduction of the bathymetric surface or its derived views (depth images, depth contours and 3D surface models) would make it practical to provide navigators with a "true present depth" Electronic Chart.
- Another possibility is to view the data as a 3D model with fixed viewing angle pointing downward and forward of the vessel's present heading. As the vessel navigates across the bathymetry, the 3D view could refresh every few seconds to reflect the vessel's current location and heading. The EC screen image seen by the navigator would simulate the view looking forward over the ship's bow. However, the formerly opaque ocean surface would be rendered perfectly transparent so the ship would appear to be flying over the seafloor. An adjustable vertical exaggeration of the model would heighten this effect. A vertical offset equal to the vessel's draft would displace the navigators viewpoint down to the keel of the vessel and an artificial water surface could be added to the graphic to provide viewing perspective.
- A third interesting viewing possibility arises from the ability to rapidly extract depth profiles from a computed bathymetric surface. Any course vector, or series of course vectors, drawn onto an EC's display of the depth image also describes a profile on the bathymetric surface model. This depth profile could be instantly extracted and drawn onto a separate window or display device. Depth under keel is the most critical dimension concerning navigators. Depth profile extraction thus permits a navigator to focus on the under-keel depth that is the most critical factor in any proposed route. By using the corresponding profile of vertical uncertainty on the stochastic surface, 95% deep and shallow probabilities for this depth profile could also be overlaid onto the graphic. Tidal predictions could also be used to compute and display true depth under keel for the present location as well as all the projected locations during the rest of the voyage.

5) Since the interpolated depth matrix is identical to the raster imagery so prevalent in today's computer graphics industry, standard image processing software can be used to render the bathymetric model more visually pleasing. The simplest image processing approach is to re-assign the color of each depth in the image using a "stair step" color LUT to effectively discard most of the vertical resolution of the depth image. This LUT could contain only those few standard colours which are used on nautical charts. The depth image re-mapped in this manner would immediately take on an appearance somewhat similar to a traditional paper chart. More sophisticated image processing algorithms such as edge finding and noise filters could operate on the full depth resolution of the depth matrix to enhance the appearance of seafloor morphology. It is important that any image enhancement operation only alter what the navigator might wish to see. All depth values in the original interpolated depth matrix must be stored permanently in the EC's database. The bathymetric surface's *spatial* integrity must be maintained for use by non-

graphical guidance algorithms that run as continual background tasks in an EC environment.

All of the above EC capabilities are feasible using easily implemented algorithms running on today's computers ... if the survey data has been geostatistically modeled.

Before leaving the subject of bathymetric data modeling, we can state three truisms concerning safe navigation:

- 1) To establish a logically efficient infrastructure for safe navigation, the vessel positioning system should provide a level of uncertainty that is approximately equal to the level of uncertainty in the map model. It is inefficient to provide a high level of accuracy in the map model if it is not matched by that of the positioning system used to exploit it (and vice versa).
- 2) An "overly-accurate" vessel positioning system can easily encourage dangerous mis-use of poorly defined map models. The advent of DGPS and Electronic Chart systems has exacerbated this problem.
- 3) Safe navigation decisions (calculated risks) can *only* be made if both the real-time vessel position uncertainty *and* the map model uncertainty are *evident* to the navigator throughout the time and space of the voyage.

QUANTIFYING DGPS POSITION UNCERTAINTY

The first half of this paper has underlined the need to statistically quantify the spatial uncertainty of the bathymetric model being used for planning as well as that of the vessel's positioning system. IHOstat effectively solves the first half of this problem. The balance of this paper deals with the other half of the safety problem: providing navigators with valid estimates of their present GPS positioning accuracy.

This safety requirement is analogous to the need for "Integrity Monitoring" which has been a much discussed requirement for aeronautical GPS applications. While ship positioning requirements are generally not as critical as those for landing aircraft, docking large ships has similar requirements. A man-made wharf face is similar to a runway in that it can be very well positioned and displayed on an Electronic Chart. A Captain attempting to berth his ship in poor visibility would obviously benefit from an accurate and reliable bird's eye view of vessel movement relative to the wharf. However, a Captain can never have *confidence* in an EC for guiding critical maneuvers unless he has confidence that the spatial uncertainty of the positioning system is being constantly evaluated to insure adequate positioning. The 2-sigma error ellipse illustrated in Figure 2 is one means of communicating this information through the EC. Position error estimates need not be continually displayed on the screen to insure positioning integrity. A limit on permissible positioning uncertainty could also be input to an EC algorithm which would monitor integrity in the background. If the length of the semi-major axis of the computed error ellipse were to exceed that value then visual and audio alarms would be

triggered to alert the Captain to the problem. Of course in order to work, this monitoring system requires that the real-time position error estimates be statistically valid at all times.

Over the last decade the Canadian Hydrographic Service has carried out R&D aimed at improving both the accuracy and reliability of GPS positioning. The main goal of this R&D was to reduce and quantify the positioning errors which contaminate hydrographic survey data. Traditionally, hydrographers have used positioning systems far more accurate than those available to general navigation. With the advent of DGPS that safety margin has been narrowed. Navigators now have access to positioning accuracy which in many cases is better than that which was available during the original hydrographic surveys. To retain as large a safety margin as possible, it is therefore incumbent on hydrographers to carry out new DGPS positioned surveys with the utmost attention to quality control. Working towards this goal, CHS and Nortech Surveys recently carried out a simple experiment designed to evaluate and hopefully improve the methodology used for estimating DGPS position errors.

EXPERIMENTAL OBJECTIVES

Real-time GPS position error estimates are vital, both for producing and using chart products. Most GPS receivers available today provide very little in the way of position quality indicators. Navigators using basic GPS equipment only have access to the GPS constellation's Geometric Dilution Of Precision (GDOP) as a scaleless index of their present positional accuracy. This is better than nothing however strong satellite geometry is no guarantee of the integrity of the GPS data actually being observed. Since it is a scaleless number, GDOP cannot possibly provide the spatial error estimates a navigator requires to safely exploit an Electronic Chart. The firmware in a few more sophisticated (expensive) GPS receivers can output spatial error estimates derived from the residuals of the positioning solution. These are the integrity indicators which are of real value to navigators.

Research into using GPS for aeronautical applications has resulted in two popular approaches to assuring adequate GPS positioning performance. These two methodologies have evolved under the general heading of "Integrity Monitoring". They are: Receiver Autonomous Integrity Monitoring (RAIM) and the GPS Integrity Channel (GIC). RAIM analyzes redundant pseudorange information to isolate faulty satellites and exclude them from the position solution. An alarm sounds when the predicted position error of the position solution exceeds the FAA tolerance for enroute air navigation (nominally 100 metres). The GIC technique employs ground monitor stations at known locations which observe GPS performance and broadcast a message to aviators if GPS errors exceed the allowable tolerance. Both of these integrity monitoring methodologies were conceived to monitor single-point positioning for enroute air navigation.

Differential corrections are now becoming widely available through a variety of data links and greatly facilitate the Integrity Monitoring task. Much of the RAIM methodology relates to detecting "faulty" satellites so they can be excluded from the position solution. RAIM is a valid QC requirement for single point positioning,

however, if differential corrections are being constantly received at the vehicle then satellite faults can be detected as jumps in the correction. Receiver malfunctions are a different problem but can be adequately addressed through hardware redundancy. By *applying* differential corrections, satellite faults can generally be both detected *and* corrected so that ranges need not be discarded from the solution. The relevance of the GIC methodology is even more affected by the widespread availability of differential corrections. GIC ground monitors are in effect differential reference stations. However instead of providing differential corrections to users, the GIC merely broadcasts the information that single point GPS is or is not providing 100 metre accuracy.

A more useful GPS Integrity Monitoring methodology would be to:

- 1) Insure that navigators use differential GPS.
- 2) Compute and use statistically valid position error estimates.

This general approach depends heavily on the on-board GPS equipment's ability to make valid real-time error estimates. That is what the experiment described below has focused on.

EXPERIMENTAL METHODOLOGY

The prime goal of the experiment was to evaluate and improve *real-time* error estimation, however, in this experiment all raw data was simply logged and *post-processed*. This may appear to be a blunder in the methodology, however it was both legitimate and desirable due to the characteristics of the GPS processing software used for the experiment. Some background discussion of this software is in order to explain this.

Over the past 8 years, the Canadian Hydrographic Service has contracted Nortech Surveys to carry out GPS R&D pertinent to its mandate. It was important to understand the inner workings of the GPS algorithms being tested, particularly in the area of Quality Control. We therefore chose to bypass the firmware solutions provided by commercial receivers and worked directly with raw GPS data. As a result of this decision, a substantial amount of generic GPS processing software was written for CHS by Nortech Surveys. CHS christened this software "Hydrostar" and used it for a variety of GPS experiments (ref. 5, 6 and 7). Hydrostar grew over time to encompass real-time DGPS processing as well as a generic data logger and coxswain's guidance display. CHS is now using this software on production surveys.

Having developed the Hydrostar software under contract to CHS, Nortech Surveys recently purchased the rights to it from the Canadian government and are now marketing it in two versions:

- HPC (Hydrostar on PC) is for use in real-time DGPS applications.
- HPM (Hydrostar Post Mission) employs the *same* algorithm as HPC but in post processing mode. HPM is useful when no differential data link is available or as an R&D tool to analyze the same data set using a variety of processing parameters. This equivalence between HPC and HPM is what permitted the post processing used in this experiment to accurately reflect the results we would have obtained in real-time.

HPC's Main Features:

- Generic real-time PC application which decodes raw GPS observations from various receivers, currently: Trimble, Ashtech, Magnavox and Novatel.
- Operates in any of three DGPS modes: Reference, Remote and Monitor.
- Displays detailed QC parameters on position residuals, satellites, clock models and phase smoothing.
- Has audio and visual alarms to alert the hydrographer to out of tolerance positioning accuracy (Integrity Monitoring).
- Displays real-time guidance graphics suitable for either waypoint navigation or following a survey grid.
- Incorporates a hard disk data logger for both GPS data and generic serial data (e.g. digital depth sounders)
 - Controls up to 6 serial ports with bi-directional assignments on each port for flexible I/O with GPS receivers, differential data-links and generic serial sensors.
 - Uses GPS time sync pulse to update drift and offset of the PC's clock thus providing accurate time stamping of all data records logged to disk (+/-0.05 ms.).
 - Optimizes hard disk head ballistics for ruggedized logging in harsh environments.
- Uses GPS phase observations to compute vertical launch motion to be applied to logged depth soundings (heave corrections).

The goal of this present experiment was to study real-time GPS error estimation. Since HPM (post-mission) shares the same positioning algorithm with HPC (real-time), we could legitimately use HPM to simulate a variety of real-time scenarios required to carry out this experiment. For example, satellite geometry was easily varied by disabling different satellites in the HPM control file then reprocessing the same data set. Data link outages were also simulated simply by decimating the data file logged at the reference site. The main reason however for

using HPM was that it permitted to observe the effects of modifying the HPC/HPM positioning algorithm by reprocessing the same data set after each software change.

Another simplification to the experimental procedure was made by using a known static baseline to simulate DGPS positioning on-board a moving survey launch. Using a known baseline simplified logistics tremendously by providing excellent reference positions at the DGPS Remote. This permitted position misclosures to be accurately observed, but was it too simplistic? Would not mounting the remote antenna on a real survey vessel produce very different DGPS position errors? This question had already been answered in 1990 and 1991 (ref. 5 and 7). Those CHS experiments were carried out using a motorized vessel movement simulator. The oscillating mast was instrumented to provide centimeter level reference positions for the moving GPS antenna. Trimble, Ashtech and Magnavox receivers were concurrently tested on this apparatus under controlled dynamics which ranged from static to quite violent. Analysis of those results showed that modern receivers track GPS signals well enough that vessel dynamics have negligible effect on position errors. In fact, when the mechanical "survey launch" was held static, position errors were actually worse due to larger multipath effects than when the antenna was in motion. Based on those findings we concluded that, for the purposes of this experiment, it was legitimate to use a static site to simulate a DGPS positioned survey vessel.

The raw GPS data was collected over a known 10 km baseline and processed using HPM. Novatel narrow-correlator C/A code receivers were used at each site. Three hours of 1 second data were logged during each of two sessions. The experimental methodology was simple:

Step 1. Using HPM, the data sets were processed in their respective modes to simulate the real-time performance of HPC:

- The Reference site generated differential corrections.
- The Remote (survey vessel) site applied the corrections to estimate DGPS positions and *also* made position error estimates.

Step 2. The Remote site's DGPS position solutions were differenced with the site's known reference coordinates to observe the true DGPS positioning errors. These true errors are then compared to the errors which had been predicted by HPM's position solution. The error estimation algorithm was then modified to try to improve agreement between what happened and what was predicted. The effectiveness of each software modification was then observed by re-processing the data.

RESULTS AND ANALYSIS

HPM processed the data using a single epoch least-squares solution for position with Kalman filtering applied to the clock model. The pseudorange were phase-smoothed using a Kalman filter algorithm recently implemented which makes use of the rate of change of differential corrections to aid phase velocity

measurements. This phase smoothing strategy is particularly helpful in counteracting the effects Selective Availability has on phase velocity. The position error estimates output by HPM were "tuned" by varying the method used to estimate covariance information about both pseudoranges and phase measurements.

The pseudo-range variance model used by HPC/HPM has two cases:

- When the solution is single point, the pseudorange variance is a combination of the broadcast URA + pseudorange noise + multipath estimate + ionospheric variance + tropospheric variance. The ionospheric and tropospheric models are functions of satellite elevation. The broadcast URA encompasses satellite ephemeris and clock errors. In the presence of Selective Availability URA is the dominant factor and results in strong favoritism towards Block I satellites in the solution.
- When differential corrections are being applied, the pseudorange variance is a combination of the factors listed above + an empirically derived "de-correlation factor" of 5% of the tropospheric correction + the variance of the differential corrections themselves.

In the first step of this paper, the contour envelopes describing the accuracy of a bathymetric surface were referenced to the 95% confidence level. This was considered to be the statistical criteria that a prudent navigator would use for evaluating spatial risk. This same 2-sigma confidence level was used to "tune" the pseudorange error model using a lengthy trial and error approach. If exactly 95% of the observed position errors were less than their predicted value, then the error estimation would be statistically perfect. Using this feedback criteria, an empirical pseudorange variance model was devised over approximately 10 iterations of software modifications. The before and after effect of those modifications is illustrated below.

Figure 3 shows HPC/HPM results for a typical one hour period of 7 satellite DGPS data. The upper two lines represent the observed position errors (solid line) and estimated position errors (dotted line) produced by HPM prior to any filter modifications. The lower two lines are the observed and predicted position errors, produced from the same data set, after modifying the range variance model. The improvement is striking! Not only is the error estimation become more statistically valid, the positioning accuracy has also greatly improved. Position errors improved from the 1.0 to 2.5 metre level when processed with the old filter to 0.1 to 0.4 metres when the same data was processed using the modified range variance model.

It is worth noting that, prior to this experiment, the HPC/HPM generic solution had been experimentally compared to position solutions computed by firmware in the 4 receivers HPC/HPM presently decodes (Trimble, Ashtech, Magnavox and Novatel). In those trials, HPM's observed positioning accuracy was consistently equal to or better than those output by the embedded firmware (ref. 5, plus in-house testing). The upper two lines in Figure 3 represent that previous level of HPM performance. The lower two lines represent the improved performance. What caused such a striking improvement in *both* positioning accuracy and the software's ability to predict that accuracy?




FIG. 3.- HPM/NovAtel Static Position Accuracy Test Using a Priori and NovAtel Measured Range Accuracies.

As incremental modifications were made to the range variance model, it was found that the observed position performance and the observed error prediction performance were directly linked ... as one improved so did the other. While many combinations of scaling factors and offsets were tried within the range variance model, one modification stood out as making the largest improvement in overall performance: we substituted the Novatel receiver's internally computed range measurement variance values in place of the constant a priori value that is typically used in least-squares adjustments.

The Novatel's range variance estimate is a by-product of its proprietary "narrow-correlator" pseudorange measurement technique (réf. 8) and is output in the raw observation message. While the Novatel receiver computes and outputs these estimates of range variance, they are not currently exploited by the position solution in its firmware. According to the receiver's documentation, this range variance estimate is generally very small and represents the thermal noise of the pseudorange measurement. It appears that this statistic represents more than just thermal noise. Even after the receiver had fully warmed up, we observed that the logged range variance estimates were generally different for each satellite tracked. While generally very small, the error estimates would occasionally vary by orders of magnitude. Each satellite tracked by the Novatel was thus generating a dynamic range error estimate which was used directly in HPC/HPM's weighting matrix. We attribute much of the improved performance to exploiting this dynamically computed covariance matrix.

Figure 3 depicts data that was collected under ideal 7 satellite coverage. To see how the new filter reacted to more challenging situations, the data was re-processed using both degraded satellite geometry and in single point positioning

mode. Figure 4 illustrates the same 1 hour period only this time satellites were removed from the solution every 8 minutes. We see that the gains in both positioning performance and error prediction performance remained very noticeable. Figure 5 illustrates the performance during intermittent single point positioning (8 minute data link "outages"). As expected, the single point errors go off scale, however, they are still well predicted and reaction time is instantaneous (a color graph is useful for discerning this). To show what is happening off the scale in Figure 5, Figure 6 illustrates real and estimated positioning errors in single point mode for the entire data set. For legibility on the black and white graph, only results for the new filter are presented on Figure 6.




FIG. 4.- HPM/NovAtel Static Position Accuracy Test
A Priori and Measured Accuracies in Reduced Coverage.

The goal here was to make error estimates valid at the 95% confidence level. A perfect performance would yield exactly 95% of the observed position errors that are less than their estimated value and 5% greater than predicted. In real-life populations, statistical perfection is impossible to attain. However, statistical biases can be detected and hopefully reduced through more realistic covariance information. Figure 3 to 6 indicate that, while the error estimates were far from statistical perfection, considerable improvement was made as a result of modifying the range variance model. In Figure 3, the old model had very poor statistical validity: only 25% of the real errors were below their predicted value ... a dangerously optimistic view to present to navigators. After tuning the filter this performance was improved to 64% ... still not the perfect 95% but a significant improvement. While not perfectly valid from a statistical perspective it's worth noting that the maximum absolute difference between real and predicted errors was still only about 0.2 metres. In the old filter these discrepancies reached about 0.8 metres.




FIG. 5.- HPM/Novatel Static Position Accuracy Test
A Priori and Measured Accuracies in Interrupted DGPS Service.

The graph in Figure 4 shows performance under degraded satellite geometry. From an error prediction point of view, this was the worst case scenario and that is reflected in the statistics for both the old and new range variance models. The 2 sigma statistical performance of the new filter was only slightly better than the old filter (56.3% vs. 51.2%). However the absolute value of the discrepancies between real and estimated must be considered. In the old filter these ranged up to 1.2 metres while using the new filter produced a maximum discrepancy of only 0.3 metres.

Figure 5 shows the effect of intermittent differential data link outages. The error estimates respond instantly to the degraded performance and the overall statistical validity is also very good. For the new filter, 83% of the real errors were less than their estimated value while for the old filter this figure was only 31%. The processing of exclusively single point data shown in Figure 6 produced even better error predictions (86.4% of observed errors were less than predicted). The worst discrepancy for single point processing run occurred at minute 52 when the observed 2D position error was 24 metres while the 2 sigma estimated error was only 15 metres.

The conditions of this simple experiment were too limited in scope for conclusive statistics to be derived. Further testing will be carried out this year using a much larger and diverse data set. This data will be used to do some additional fine tuning to HPC/HPM's filter. All of the modifications done to date were made to the pseudorange variance model. If more extensive testing continues to show optimistic biases in the error estimations then an empirical rule set of offsets and/or scaling factors will be determined and applied directly to the position error estimates.




FIG. 6.- HPM/Novatel Static Position Accuracy Test
Single Point Statistics Using NovAtel measured Range Accuracies.

Regardless of future improvements, when viewed from an absolute perspective (i.e. the magnitude of the difference between real and estimated position errors) HPC/HPM's present error estimation performance in differential GPS mode was *always* better than a metre and generally less than a few decimetres. This level of error estimation should be sufficient to provide Integrity Monitoring for even the most stringent navigational task. It is definitely sufficient to support any navigation task involving marine Electronic Chart.

CONCLUSION

Spatial uncertainty is inherent to the approximate bathymetric models navigators use to plan their voyage. Spatial uncertainty is also inherent to the GPS positioning system they use to follow their chosen route. Both uncertainties must be evaluated and considered in order to make safe navigational decisions in an Electronic Chart environment. IHostat provides a standardized and statistically optimal method of evaluating the spatial errors inherent to a bathymetric map model. The bathymetric and stochastic surfaces computed by the program also possess some very useful visualization characteristics for EC applications. Real-time GPS position error estimation capability has been demonstrated by HPC that is well within the Integrity Monitoring requirements of vessel navigation. When displayed graphically in an Electronic Chart environment, the IHostat and HPC error estimates greatly augment a navigator's intuitive ability to assess the risk of planned

maneuvers. This statistically derived spatial information permits navigators to practice their art with greater confidence.

Bibliography

- [1] KIELLAND P., BURROWS K., WARD B., DAGBERT M., VELBERG R., (1993) "Towards IHOstat: IHO approved software that evaluates the quality of bathymetric data", Lighthouse No. 47-Spring 1993, pp. 19-32.
- [2] KIELLAND P., and DAGBERT M., (1992), "The use of spatial statistics in hydrography", International Hydrographic Review, Monaco, LXIX(1), pp. 71-92.
- [3] CHILES J.P. and CHAUVET P., (1975): "Kriging: a method for cartography of the sea floor", International Hydrographic Review, pp. 25-41.
- [4] DAGBERT M., KIELLAND P., L'ESPERANCE M. and COWAN A. (1992): "Geostatistics to assist hydrographic survey design", Proceedings of Fourth International Geostatistical Conference, Troia, Portugal, Sept. 1992.
- [5] KIELLAND P. and NEUFELDT D. (1991) "A performance evaluation of three state of the art GPS receivers for hydrographic surveying", Proceedings of the Fourth International Technical Meeting of the Satellite Division of the Institute of Navigation, Albuquerque N.M Sept. 11-13, 1991.
- [6] LACHAPELLE G., W. FALKENBERG, M. CASEY and P. KIELLAND (1989), "Hydrostar: a computer program for accurate GPS real-time kinematic differential positioning", Lighthouse No. 39.
- [7] LACHAPELLE G., W. FALKENBERG, D. NEUFELDT and P. KIELLAND (1990), "Marine DGPS using code and carrier in a multipath environment", Lighthouse No. 41, pp. 33-37.
- [8] VAN DIERENDONCK A.J., FENTON P. and FORD T. (1992), "Theory and performance of narrow correlator spacing in a GPS receiver", Navigation, Vol. 39, No. 3, pp. 265-283.