Enhanced Total Propagated Uncertainty (TPU) in CARIS HIPS and SIPS

Burns Foster[1] (presenter), Karen Hart[2], Grant Froelich[3], Bill Lamey[1]

1. CARIS, 115 Waggoners Lane, Fredericton, NB CANADA E3B 2L4, 506-458-8533
2. CARIS USA, 415 N. Alfred St., Alexandria VA, 22314, 703-299-9712
3. NOAA National Ocean Service, Office of Coast Survey, Hydrographic Surveys Division, Pacific Hydrographic Branch, N/CS34, 7600 Sand Point Way N.E., Seattle, WA 98115, 206-526-4374

Abstract

The Total Propagated Uncertainty (TPU) - the computed horizontal and vertical uncertainty associated with each sounding – helps identify and remove bias from processing, however there has traditionally been difficulty analyzing and troubleshooting its individual uncertainty components. Based on a set of requirements provided by personnel at the National Oceanic and Atmospheric Administration (NOAA), CARIS[1] retooled the compute TPU dialogue and created a TPU analysis tool in HIPS[2] and SIPS™ to more efficiently solve this problem.

NOAA needed finer control over the TPU algorithm in HIPS to allow application of a portion of the a-priori modeled uncertainty over individual components of real-time uncertainty in instances where the real-time uncertainty could be suspect. Therefore, the Compute TPU dialogue now allows users to control individual uncertainty sources - static source, model, or real-time source - for each component of the TPU.

NOAA also required a visualization tool to quickly identify specific components of the uncertainty that might be the highest contributors of uncertainty, especially in cases where the uncertainty was outside the IHO S-44 Order for the survey. CARIS developed a tool to visualize the results of the Compute TPU process and show the user the breakdown of the error contributions for a subset of soundings or an individual sounding. This includes the overall component breakdown for both vertical and horizontal error presented to the user numerically and as a percentage.

These new tools provide full control and intuitive visualization over the error model, a critical aspect for all hydrographic surveys.

[1] This term is a trademark of CARIS (Universal Systems Ltd.), Reg. USPTO.

[2] This term is a trademark of CARIS (Universal Systems Ltd.), Reg. USPTO.

Introduction

The use of Total Propagated Uncertainty (TPU) as a measure of bathymetric survey quality has become a prevalent tool in the surveyor's toolkit as a means to quantify the quality and accuracy of acquired data. However, the model used to produce a single horizontal and vertical accuracy estimate for a sounding measurement has traditionally remained opaque to the end-user, limited to researchers and hobbyist mathematicians. Even among this group, breaking down the individual components which produced this final quantification can be difficult for more than a single set of measurements. When the model produces uncertainties well outside the expected range, it can be a daunting task to open up the model and locate the largest contributors.

Personnel at the National Oceanic and Atmospheric Administration (NOAA), Hydrographic Surveys Division (HSD), approached CARIS to define a new set of tools to aid them in situations where field conditions, data corruption, user error, or other problems appear as unacceptable uncertainty estimates in their field surveys. From a practical standpoint, this involved modification to the existing TPU algorithm in CARIS HIPS and SIPS, and the introduction of a new analysis tool.

TPU as a Modeling Criteria

The intention of TPU is an assessment of quality of the recorded data. When creating data products from the source soundings, methods exist which use the TPU values as part of the modeling technique. A common modeling technique in processing multibeam surveys is using the Combined Uncertainty and Bathymetric Estimator (CUBE) algorithm, developed by Dr. Brian Calder at the Center for Coastal and Ocean Mapping (CCOM), University of New Hampshire. This algorithm, developed in conjunction with a TPU library, uses the TPU values in calculating the most likely depth value for a given model cell. In addition, the criteria for a sounding contributing to a given cell is based on IHO Survey Order criteria (International Hydrographic Organization, 2008). A cut-off limit is defined for the given depth range, calculated by:

$$\pm\sqrt{a^2 + (b^2 \times d^2)}$$

Where:

a represents that portion of the uncertainty that does not vary with depth

b is a coefficient which represents that portion of the uncertainty that varies with depth

d is the depth

$(b^2 \times d^2)$ represents that portion of the uncertainty that varies with depth

The values for this calculation are sourced from the IHO S-44 Survey Order specifications:

S-44 Standards for horizontal and depth accuracies.

Order	Special	1a	1b	2
Description of areas.	Areas where under-keel clearance is critical	Areas shallower than 100 metres where under-keel clearance is less critical but *features* of concern to surface shipping may exist.	Areas shallower than 100 metres where under-keel clearance is not considered to be an issue for the type of surface shipping expected to transit the area.	Areas generally deeper than 100 metres where a general description of the sea floor is considered adequate.
Maximum allowable THU 95% *Confidence level*	2 metres	5 metres + 5% of depth	5 metres + 5% of depth	20 metres + 10% of depth
Maximum allowable TVU 95% *Confidence level*	a = 0.25 metre b = 0.0075	a = 0.5 metre b = 0.013	a = 0.5 metre b = 0.013	a = 1.0 metre b = 0.023

Table 1: IHO S-44 Survey Order Specifications (simplified)

Based on the survey order under which a survey is undertaken, the user must be aware of the uncertainty criteria required for the acquired data.

Compute Total Propagated Uncertainty

Calculating uncertainty estimates for soundings using the Compute TPU function has been a tool in HIPS and SIPS since 2004. Based on the TPU library provided by CCOM, it was initially implemented as a static model that users can manually input component error estimates to be used during the computation for each sounding. This library has been extended over the years for new systems and configurations. For example, recently, hardware vendors have begun providing real-time uncertainty estimates for each sample for various sensors, calculated and stored during acquisition or in post-processing. Uncertainty RMS values computed for post processed navigation and attitude information to create a smoothed best estimate of trajectory (SBET) can also be used in the TPU computation. The TPU library has been extended to ingest these real-time estimates and use these values in place of the static model calculation where appropriate.

In previous versions of the software, the user could only choose between “real-time” and “vessel” (static) for their source uncertainties, which defined which type was favoured over the other. Choosing "real-time" would use any real-time estimates, like RMS or sonar uncertainty, in place of the static model. If the “real-time” source data was unavailable when “real-time” was chosen, the software would fall back to the static “vessel” model for computing the uncertainty component.

The limitations of this approach appear when a real-time uncertainty component is simply unavailable for one or more sensors, it contains corrupt or miscalculated information, or a user wants to purposely use the static model for certain components (even if real-time uncertainty is available), for testing and comparison purposes for example. Previous versions of the software did not provide the option to pick and choose, on a sensor-level basis, which uncertainty source to use. The static settings would only use the static model defined by the user for all sensors,

while the "real-time" settings would use any real-time uncertainty loaded for all sensors if available, regardless of any potential problems.

This lead to the first modification to the software, which involved providing a list of sensor types and allowing the user to customize which uncertainty source would be used when computing TPU (Figure 1).

Figure 1: Compute TPU Dialogue

The dialogue provides an all-static and all-real-time option, in addition to the per-sensor customization. During processing, the user can apply this configuration on a per-line basis or group of lines as required, with the configuration used being logged on a per-line basis.

TPU Visualization

With the uncertainty sources configured and the uncertainty values calculated, the larger challenge in an investigative task is visualization of the results and interrogating where required. There are a number of tools available already, by visualizing TPU on a per-sounding basis in the Subset Editor (Figure 2), or the propagated uncertainty which results from a gridding operation through an algorithm such as CUBE (Figure 3).

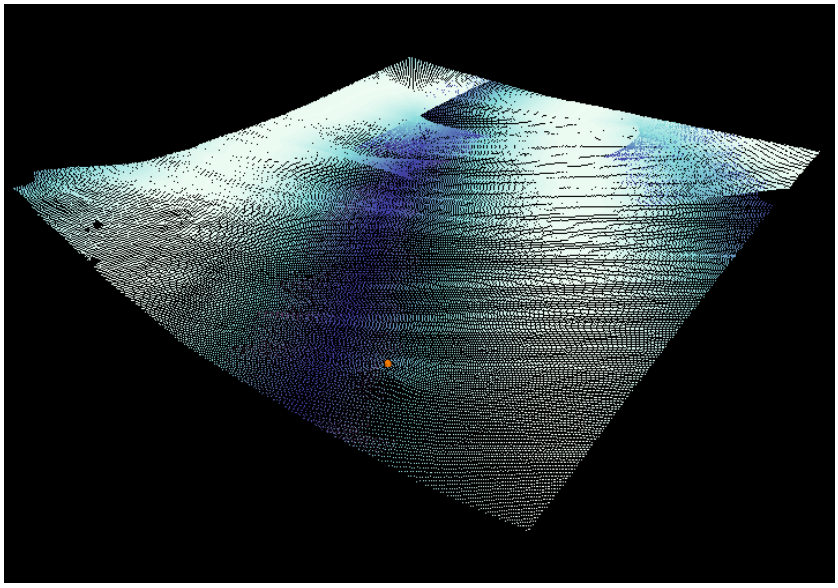

Figure 2: Adjacent Lines Coloured by Depth TPU. White is low, blue is high.

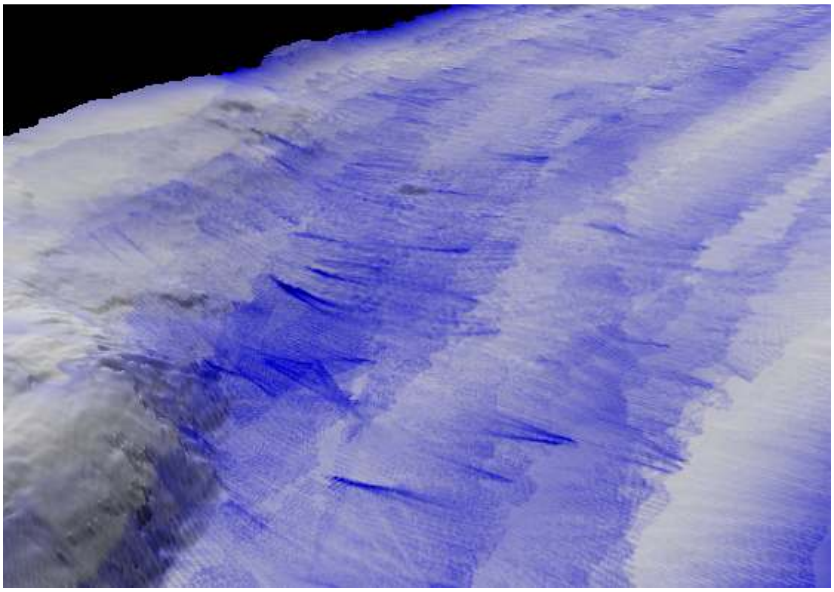

Figure 3: Gridded Representation, Coloured by Uncertainty. White is low, blue is high.

While useful in assessing the final quality of the survey data, should outliers present themselves in the computed TPU, it becomes difficult to assess what contributed to the sum of the depth and horizontal uncertainty.

The TPU Analysis tool was introduced in HIPS and SIPS version 8.1.4. Using the TPU library embedded in the application, a real-time breakdown of the uncertainty contributions from each sensor type is displayed from any bathymetry selected by the user with the tool active, specifically in the Swath Editor and Subset Editor. The TPU breakdown can be displayed as a pie chart, bar chart, or scatter plot. When multiple soundings are selected, an average of each

component is calculated and presented in the pie and bar chart displays. The tool becomes particularly useful when selecting a particular ping, or particular beam number.

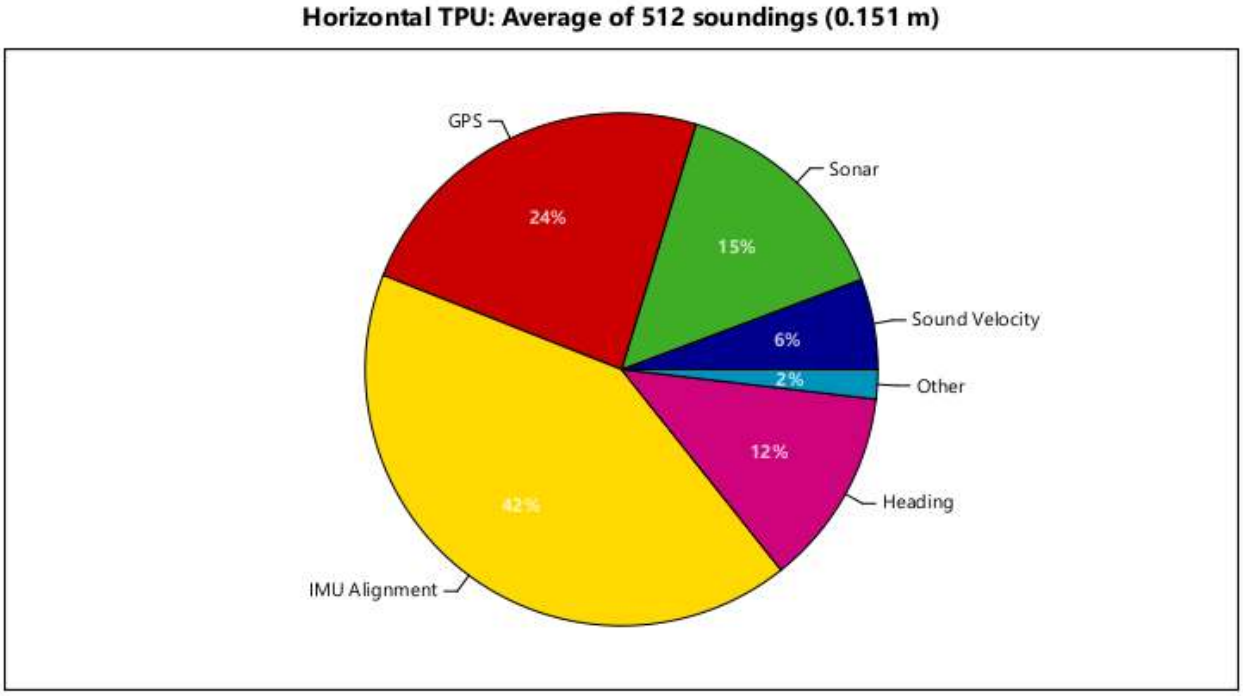

Figure 4: Breakdown of Horizontal TPU % for 1 Ping, Pie Chart

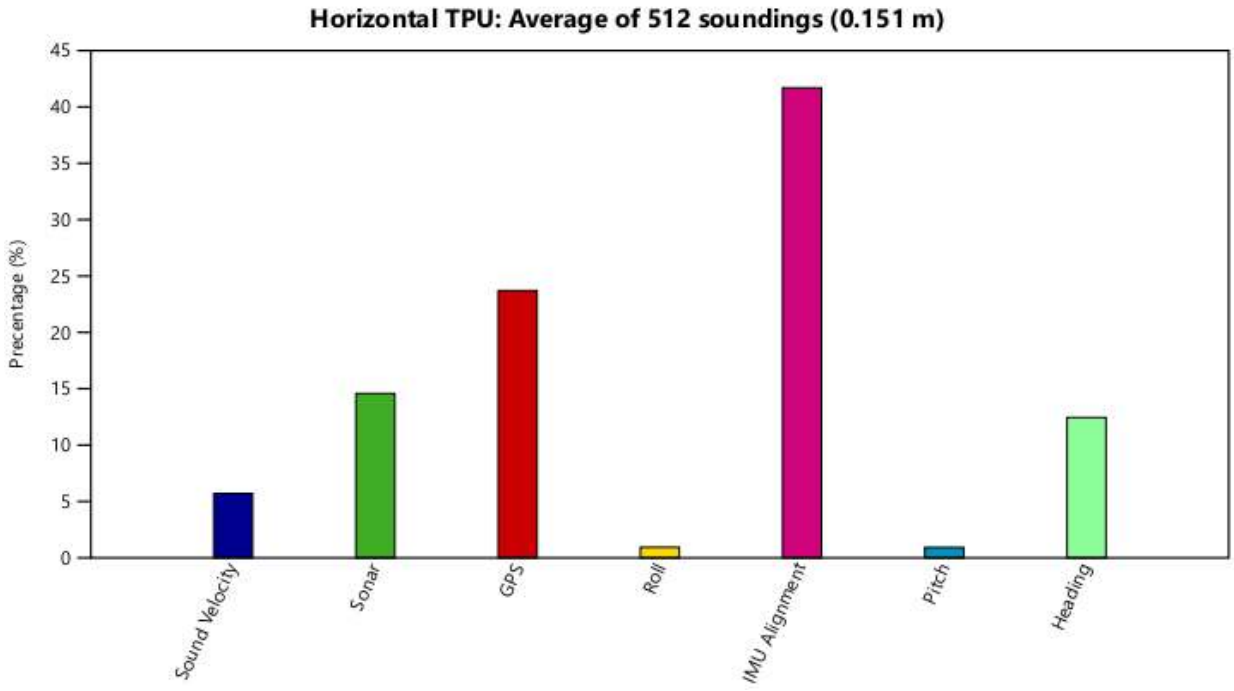

Figure 5: Breakdown of Horizontal TPU % for 1 Ping, Bar Chart

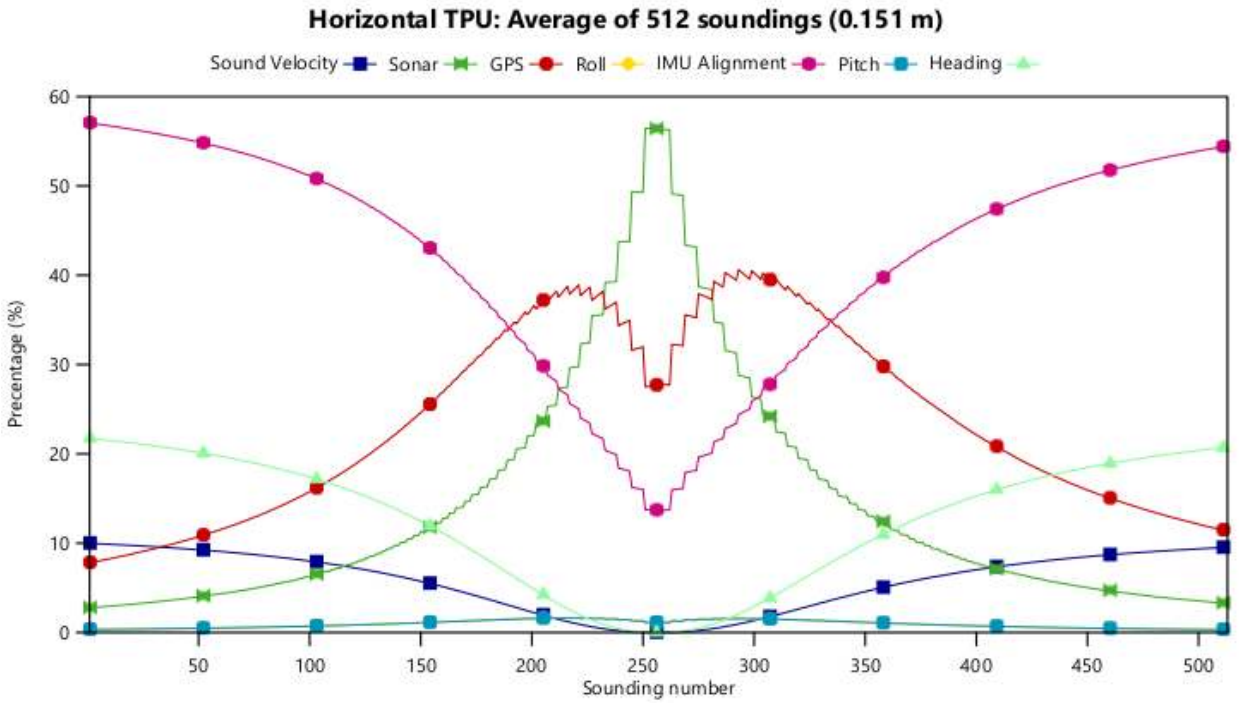

Figure 6: Breakdown of Horizontal TPU % for 1 Ping, Scatter Plot

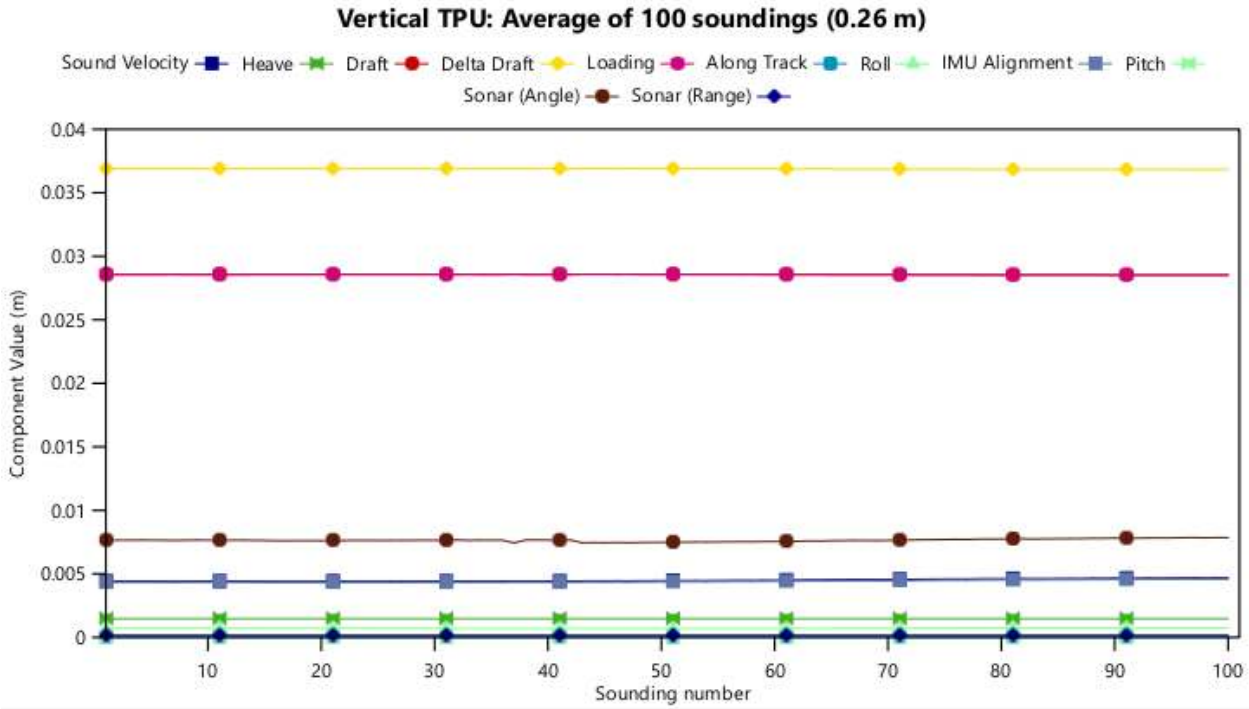

Figure 7: Breakdown of Vertical TPU values for 1 Beam, 100 Pings Along-Track, Scatter Plot

Individual uncertainty components can be turned on or off. Values can be displayed as percentages of the total TPU or absolute value.

TPU Analysis

After identifying any areas of unacceptable uncertainty, via the visualization of an uncertainty surface for example (Figure 3), the user will likely want to analyze the uncertainty of a selection of soundings. The pie chart above (Figure 4) is an example of how to quantitatively visualize how each component of the TPU for a selection of soundings relates for both horizontal and vertical uncertainty.

Unacceptable uncertainty can arise from a variety of factors, not the least of which is user error. Values for the static TPU model are entered manually by the user, sourced from manufacturer specifications, calibration routines or field experience, and every user input is a source of potential problems, the so-called "fat-finger" mistakes.

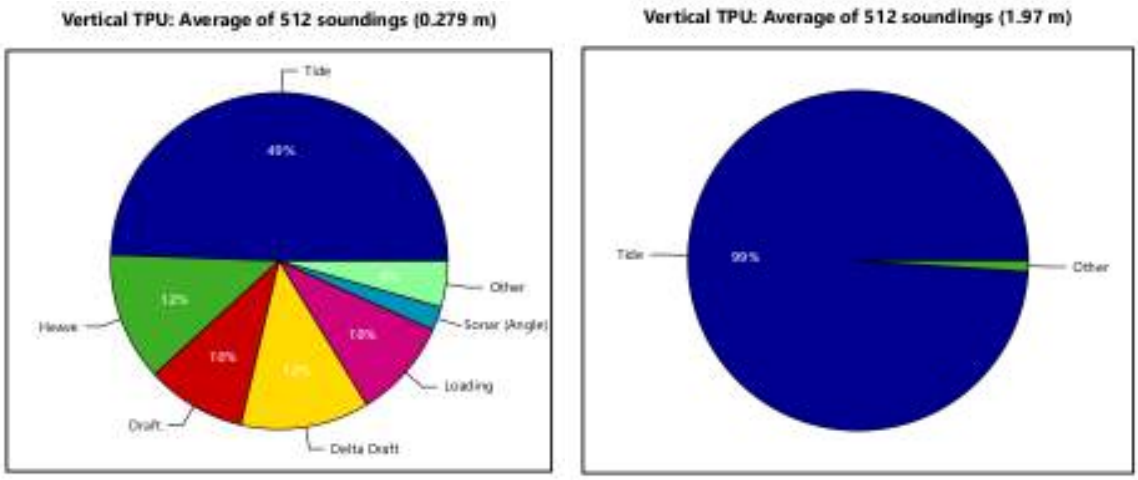

Figure 8: Vertical TPU from Static Source with Tide Blunder

Figure 8 shows a normal (left) Vertical TPU breakdown versus a "fat-finger" mistake (right) in entering the Tide uncertainty. The first indication would be a high value in the resulting products (from the *~2m* uncertainty), which can then be traced to the disproportionately high Tide contribution. A review of the processing logs confirms a value entry of '*1.0m*' instead of the more appropriate '*0.1m*' value.

Real-time uncertainty estimates can also be examined. As they are computed directly during acquisition, or on a per-sample basis in post-processing, they can be a more realistic estimate of uncertainty, particularly against a static model using manufacturer-stated accuracies. However, given the nature of real-time data streams they are more susceptible to noise, interference and drop-outs. When this happens, the static models can be used in place of the real-time values.

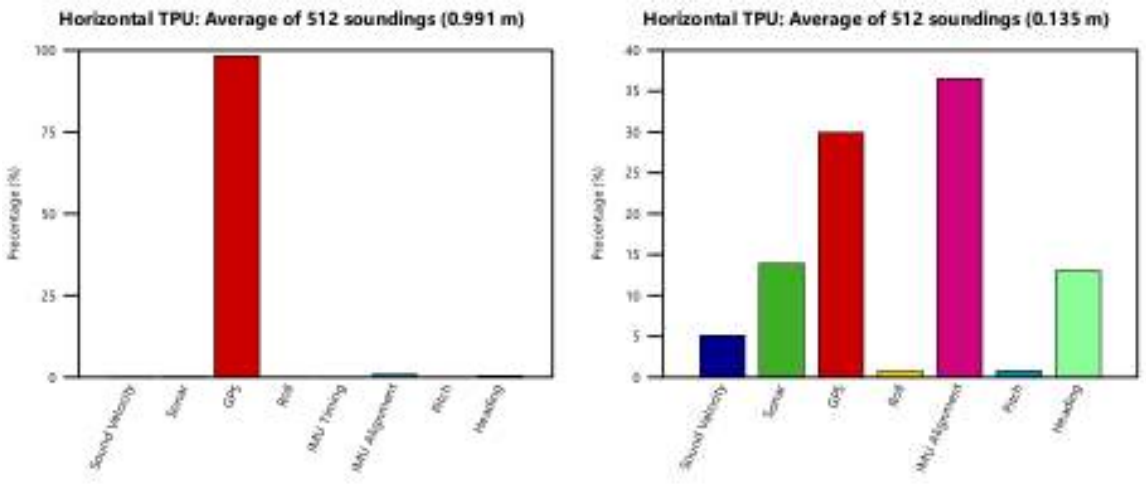

Figure 9: Static (left) versus Real-time (right) Horizontal TPU

Figure 9 displays the difference between static and real-time horizontal models for a single ping. The positioning component (GPS) represents a significant contributor in this regime versus the other components, and overall models a much larger uncertainty value (*0.991m* average for the ping) against the real-time value (*0.135m* average for the same ping). In fact, such a difference might warrant an investigation into the source values of both the static and real-time models. In this case, the static value is calculated using a generalized value of unknown origin at *0.5m* (1 sigma), which is a very pessimistic number for modern positioning systems under normal operating conditions.

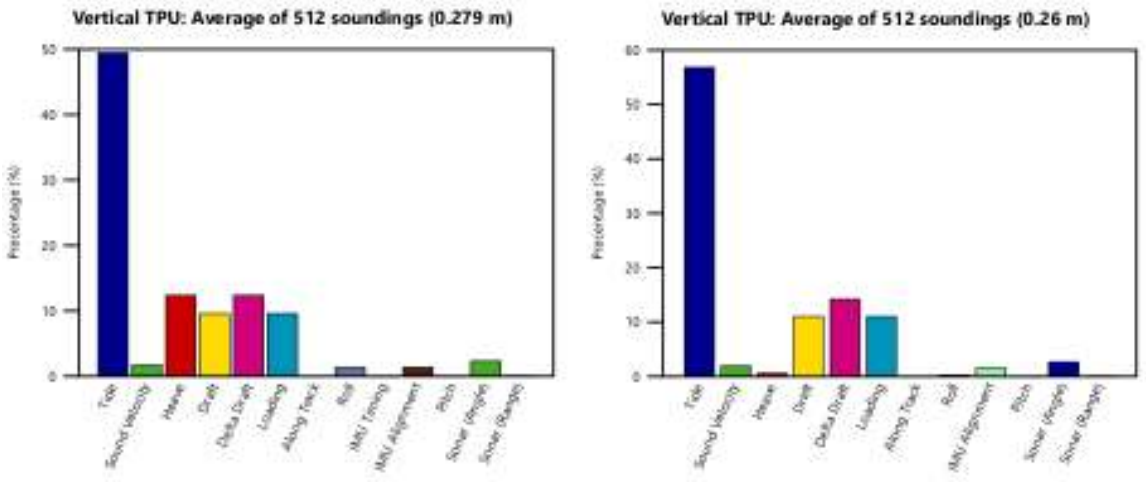

Figure 10: Static (left) versus Real-time (right) Vertical TPU

Figure 10 displays the difference between static and real-time vertical models for a single ping. While there are differences in various components (Heave being the most significant), overall the two models are generally compatible. In this case, static heave uses the static model value of 5cm +/- 5% of heave, whereas the real-time heave is post-processed delayed heave, and as such has a much smaller uncertainty.

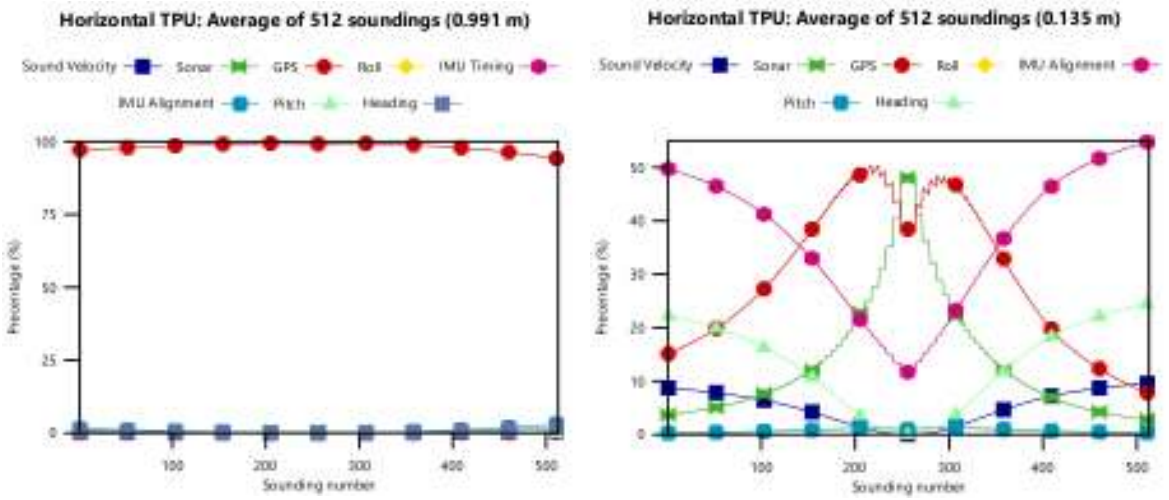

Figure 11: Static (left) versus Real-time (right) Horizontal TPU, Scatter Plot

Figure 11 shows how a single ping can be analyzed using the scatter plot, in this case again showing the static and real-time models. As seen previously, the GPS component of the static model is the largest contributor, and in this case obscures the other values. Individual components can be hidden, and the GPS component is hidden in Figure 12.

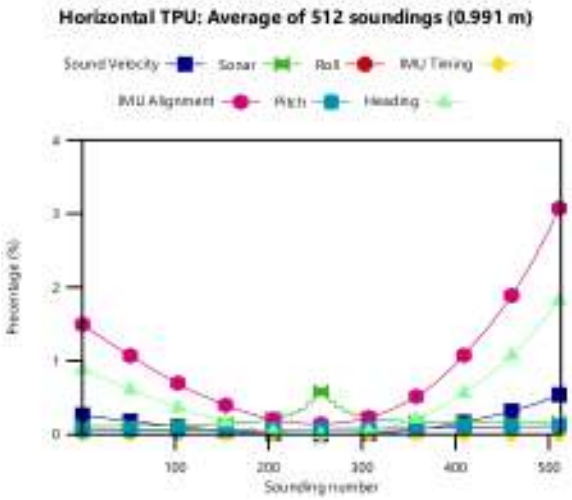

Figure 12: Horizontal TPU, Scatter Plot, GPS Component Disabled

With the GPS component hidden, a comparison against the real-time model in Figure 11 shows the two models are partially compatible. However, some differences can be seen from nadir to the swath edges in IMU alignment and heading. This may indicate differences between the standard TPU model and the manufacturer's real-time estimate for these values, and warrants further investigation.

Line vs. Area-Based Analysis

The TPU Analysis window operates on any arbitrary selection of soundings, within both the line-based Swath Editor and the area-based Subset Editor. This lends to differing types of analysis depending on each context. TPU Analysis within the Swath Editor is generally more useful in interrogating TPU data within a single line, on individual pings or specific beam numbers on that line. Conversely, analysis within the Subset Editor is generally more useful when uncertainty trends are seen on a more regional level, for instance areas of high variability in the water column (sound speed), or specific unique tidal regimes. This approach is also useful when areas of high uncertainty are identified in a gridded product, where a gridding method (e.g. CUBE) is based on the IHO Survey Order criteria.

Use Case

In 2014 a NOAA contractor submitted a survey which was submitted under IHO Order 1 survey quality specifications. The acceptance procedures at NOAA include generating a layer of calculated Total Vertical Uncertainty, or "TVU-ness", on the grid (Figure 13), which is simply a difference of the per-node uncertainty against the IHO Order 1 criteria for the depth at that same node. Nodes which result in a negative value exceed allowable IHO Order 1 TVU criteria.

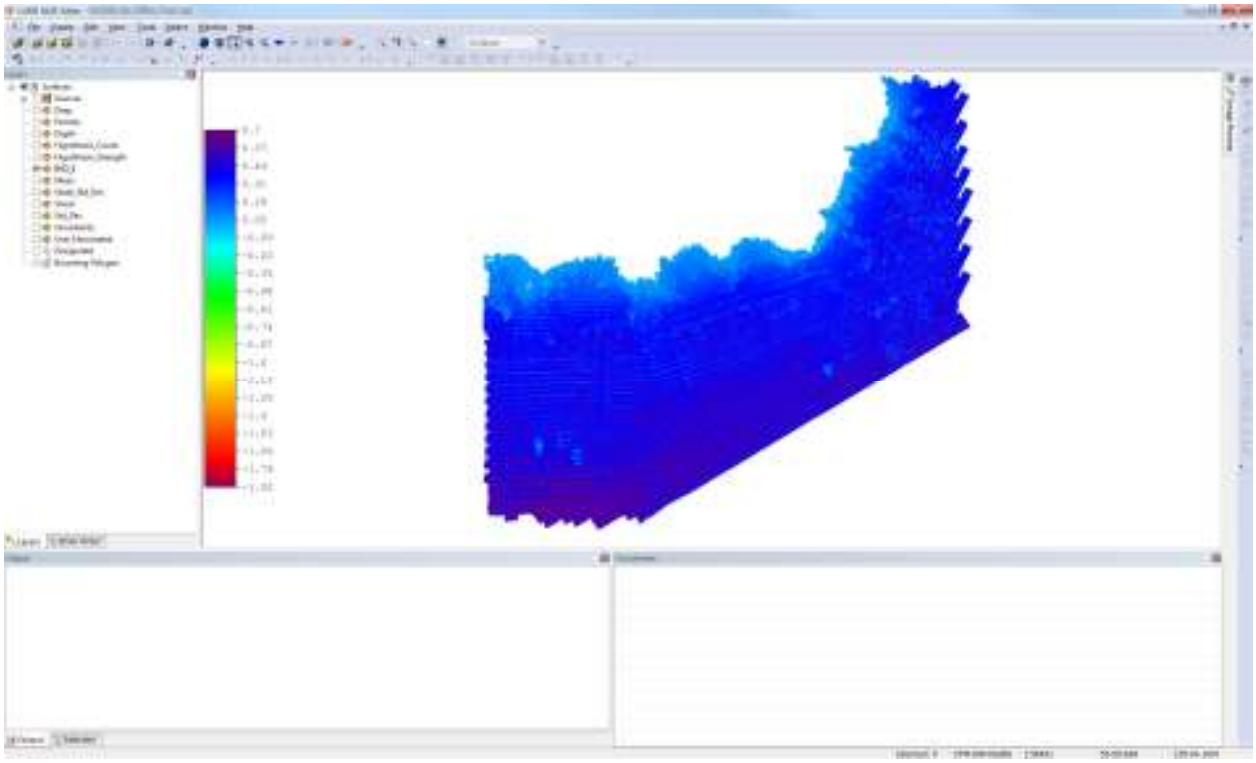

Figure 13: "TVU-ness" Layer of a Survey. Negative values are less than IHO Order 1

On inspection, the data exhibited higher than typically expected uncertainty values for the collection of sensors utilized and the operating environment of the survey. Negative values were

also discovered, meaning the propagated uncertainty of some grid nodes did not meet Order 1 criteria and failed to meet the NOAA specifications for the survey.

This prompted a thorough review of the source data's TPU components to determine the source of the higher than expected uncertainty values and to determine what steps, if any, could be taken to rectify the high uncertainty.

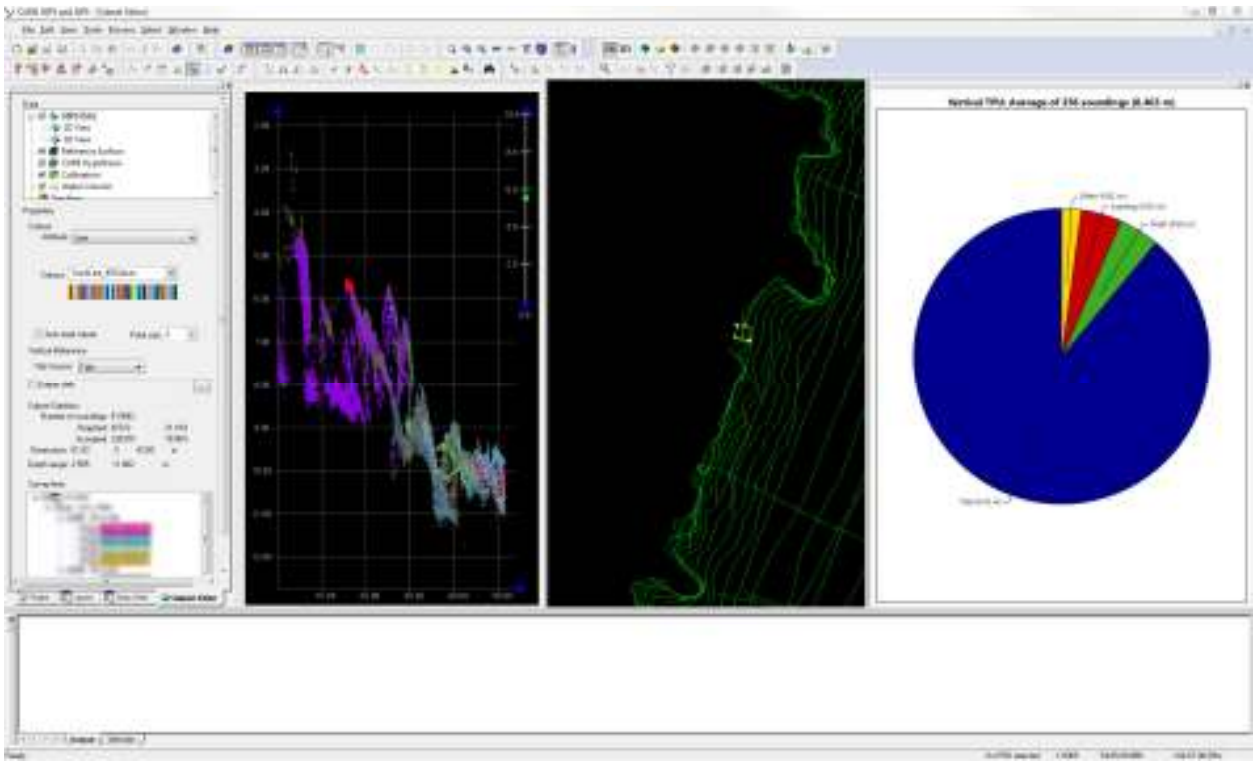

Figure 14: TPU Analysis of Survey

Figure 14 shows the TPU Analysis result from soundings selected on a shoal, specifically the Vertical TPU breakdown. The largest contributor in this sounding group by a large margin is Tide. Based on previous experience, the abnormally large tide contribution to the TVU of the soundings prompted a review of the source parameters used in their TPU computation. In this example, because the TPU Analysis tool pinpointed the problem as originating with a portion of the tides component, the Process Log, which contains the applied sound speed and tide uncertainty values was utilized as part of the investigation. The investigation of the logs revealed that the Tide Measured uncertainty was accidentally entered at a 2-sigma level instead of the required 1-sigma confidence and had subsequently doubled the Tide TVU component, leading to the higher than expected TVU values.

Prior to the introduction of the TPU Analysis tools, if the questionable uncertainty value did not lie with the applied sound speed or tide uncertainty values recorded in the Process Log, the investigation would end without a definitive conclusion, the erroneous TPU values could not be repaired, and the survey data would be treated as suspect. Now, whatever the source of the

questionable uncertainty, an in-depth investigation using the TPU Analysis tools will reveal the source and allow remediation to take place.

Conclusion

By visualizing the uncertainty components, which are ultimately used to quantify the quality of the end product, some confidence can be gained that the values put into the uncertainty model are free of error and properly modeled. In cases where products show areas of unacceptable uncertainty, this new tool allows a precise interrogation and break-down of the various contributors. Users can then establish if a re-survey is required, and if so, identify a focused area to develop further, ultimately saving time, money and frustration.

References

International Hydrographic Organization, 2008, *IHO Standards for Hydrographic Surveys, Special Publication No. 44,* 5th Edition, International Hydrographic Bureau, Monaco, 36p.

Author biographies

Burns Foster is the HIPS and SIPS Product Manager at CARIS headquarters and has been with the company for 6 years. His first 4 years at CARIS were as a Technical Support Consultant, providing training and consultation in various CARIS products. He holds a Bachelor of Science in Engineering degree (Geomatics) from the University of New Brunswick.

Karen Hart is a Senior Hydrographic Consultant at CARIS USA and has been with the company for nearly 7 years. Prior to joining CARIS, she worked as a Marine Scientist/Hydrographer for 6.5 years at SAIC. She has an MS degree in Geological Oceanography and a BS degree in Geology.

Grant Froelich is a senior Physical Scientist with NOAA's Pacific Hydrographic Branch in Seattle, WA. He has worked at NOAA since 2002. His current tasks at NOAA include completing quality assessment reviews of hydrographic surveys, generating chart update products, assisting and performing quality assurance while deployed to hydrographic vessels, and serving on the NOAA Field Procedures Manual Technical Review Board. He is a certified as a Hydrographer by the American Congress on Surveying and Mapping/ Hydrographic Society of America and holds a BA (2001) in Marine Science from the University of San Diego.

Bill Lamey is the Research and Development Manager at CARIS' headquarters located in Fredericton, New Brunswick, Canada. He has been with CARIS for over 15 years and has been a leader in the development of Hydrographic and Bathymetric COTS solutions provided by CARIS. He holds a Bachelor of Computer Science from the University of New Brunswick.