

What Can Variable Resolution Do For You?

Authors

Burns Foster – Product Manager, HIPS and SIPS

burns.foster@teledyne.com

Teledyne CARIS, 115 Waggoners Lane, Fredericton NB E3B 2L4, (506) 458-8533

Karen Hart – Director, CARIS USA

karen.hart@teledyne.com

Teledyne CARIS USA, 415 N Alfred Street, Alexandria VA 22314, (703) 299-9712

Tami Beduhn – Account Manager, CARIS USA

Tami.beduhn@teledyne.com

Teledyne CARIS USA, 415 N Alfred Street, Alexandria VA 22314, (703) 299-9712

Abstract

Variable Resolution technology is now available in a commercial product with the release of CARIS¹ HIPS² and SIPS³ 10.0. This new technology is intended to solve one of the most significant limitations in processing today, where multiple gridded products must be maintained when a survey area covers a large range of depths, and/or is surveyed using multiple platforms producing varying data densities.

This paper will explore the workflow changes involved in the use of variable resolution, and will quantify efficiency gains through the use of a single model for product generation. We will also touch on impacts to industry standards and deliverables, such as those under IHO S-44.

¹ This term is a trademark of Teledyne CARIS, Reg. USPTO and CIPO

² This term is a trademark of Teledyne CARIS, Reg. USPTO

³ This term is a trademark of Teledyne CARIS, Reg. USPTO

Table of Contents

Contents

Authors.....	1
Abstract.....	1
Table of Contents.....	2
Introduction.....	3
Methods.....	4
Processing Workflow.....	4
Scenario 1.....	4
Results.....	5
Discussion.....	8
Scenario 2.....	10
Results.....	10
Discussion.....	11
Scenario 3.....	13
Results.....	14
Discussion.....	17
Specifications and Deliverables.....	18
Conclusion.....	21
References.....	22
List of Figures.....	22
List of Tables.....	22
Trademarks owned by Teledyne CARIS, Inc.	23
Copyright owned by Teledyne CARIS, Inc.	23

Introduction

One of the most significant challenges in processing high-density remote sensing data, specifically from multibeam sonar or LiDAR (Light Detection and Ranging), is creating detailed and efficient computed models that can be used for data analysis and product generation. The nature of remotely sensed data means the density of these data varies spatially based on several factors, e.g. range to target, scanning patterns, limitations of the transmission medium, etc. For example, acquired data from traditional multibeam surveys from surface vessels becomes sparser as depth increases (i.e. range to target); while for a bathymetric LiDAR system, the range limitation is compounded by the limited transmissibility of light through water. For efficient data collection, many surveyors are increasingly using a combination of sensors and platforms to survey a given region in a cost-effective and efficient manner. Yet this adds complexity to data processing when attempting to create a single contiguous and homogeneous surface from this combination of data sources. For this and other scenarios, the fundamental difficulty is the spatial variability in the scanned data density. Standard gridded models require a single grid spacing across a given region of interest, forcing the operator to compromise between sufficient coverage in all areas (e.g. choosing a coarse resolution to account for areas of sparse data) and maintaining a spacing fine enough to pick out features of interest (e.g. keeping a resolution fine enough to meet IHO feature detection requirements for a given depth range).

Variable Resolution technology supports the creation of a single gridded product from multiple sensor sources and/or across large regions even if the data density varies by increasing depth (e.g. multibeam sonar data) or other patterns. The source data distribution is automatically analyzed to determine an optimum resolution (grid spacing) for each region of the gridded surface. The grid spacing is allowed to vary by region in an area of a given dataset, yet continuity at the boundaries between each of these regions is maintained. Variable resolution technology, therefore, simplifies the creation and maintenance of multiple gridded products down to one, increasing processing efficiency and ease of data management, while concurrently determining an optimum resolution that preserves feature detection capabilities and provides sufficient data coverage.

Methods

Processing Workflow

A traditional processing workflow would generally begin with importing the raw data, applying correctors such as tides and post-processed inertial solutions, and computing geo-referenced positions for each point. At this stage, data cleaning and quality control begins, ideally by first computing a model of the dataset to assess data distribution, coverage, and other factors. Depending on the complexity of the source data, this could involve generating multiple models at differing resolutions to capture all areas of the dataset. A Variable Resolution grid will reduce the number of models required to one, regardless of data complexity and extents. From this single model, data cleaning and quality control can take place and any number of products can be generated in exactly the same manner as with fixed-resolution models.

To quantify the use of Variable Resolution for data processing we will consider three scenarios:

1. a single-day inner harbour survey,
2. a larger multi-platform survey over a moderate distribution of depths and data densities, and
3. a mixed multibeam and LiDAR dataset over a coastal area.

We will assume raw data processing and QC take an equal amount of time, regardless of the modeling method, and disregard this contribution for this exercise. Additionally, while the three scenarios will focus on processing HIPS and SIPS, it is worth noting the same Variable Resolution process is available in BASE Editor.

Scenario 1

The dataset of interest for Scenario 1 is an inner harbour area, specifically the inner harbour in Plymouth, UK from the Shallow Survey 2015 Common Dataset.

- Sonar: Reson 7125
- 203,363,840 soundings
- 101 lines
- 1.04 km^2 , 48.5 *LKM*
- Minimum depth: 0.481 *m*
- Maximum depth: 41.281 *m*




Figure 1 - Plymouth, UK, SS2015

The dataset is provided as a fully processed HIPS project in the Common Dataset delivery, and as stated earlier, we will disregard processing time prior to gridding as being consistent regardless of modeling method. For this exercise we will examine the creation of a single gridded product and two Variable Resolution methods for resolution estimation, both using CUBE (Combined Uncertainty and Bathymetry Estimator) for the gridding method. Contours were generated at intervals 0, 10, 20, 30, and 40m, and the Sounding Selection method was done using a radial spacing of 5.0m on the ground.

Results

The timing results are as follows:

Method	Resolution	Creation Time	Contours	Sounding Selection	Total Time
Single Resolution - CUBE algorithm	0.5m	00:07:15	00:00:18	00:00:20	00:07:53
Variable Resolution					
- CARIS Density	Min 0.09m	00:29:17	00:07:53	00:00:41	01:00:23
- CUBE algorithm	Max 3.52m	00:22:32			
Variable Resolution					
- Ranges estimation		00:01:49	00:00:29	00:00:08	00:11:48
- CUBE algorithm	0.5m, 1m	00:09:22			

Table 1 - Scenario 1, Surface Timings

Density Estimation

The CARIS Density method with CUBE is an order of magnitude slower than traditional fixed-resolution gridding (~1 hour vs ~8 minutes). However, the density analysis method is gridding at a significantly finer resolution (as low as 9cm) than the single resolution surface at 0.5m node spacing. Figure 2 shows the resolution distribution in 0.1 m bins; in

this case we can see the dataset was gridded predominantly at a 20cm spacing, under density estimation. This much finer spacing significantly drives up the processing time for the final product.

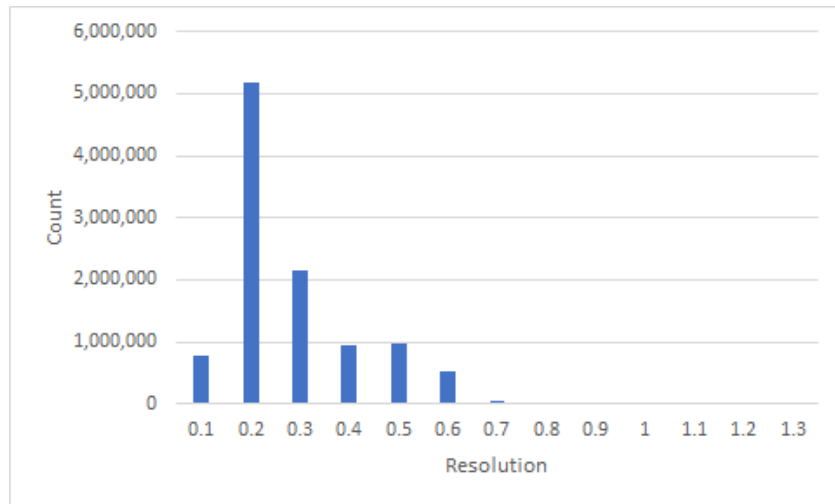


Figure 2 - Resolution Distribution in a Variable Resolution Surface

Knowing something about the data distribution, we can constrain the Density analysis to a smaller expected resolution range, which in turn reduces the total processing time:

Method	Resolution	Creation Time	Contours	Sounding Selection	Total Time
- CARIS Density	Min 0.5m	00:09:05	00:02:41	00:00:23	00:25:48
- CUBE algorithm	Max 4.3m	00:13:39			

Table 2 - Scenario 1, Additional Density Analysis

This processing time is still much longer than the fixed resolution case, but there is an important caveat to consider; at a fixed 0.5m grid spacing the following point distribution per node is observed:

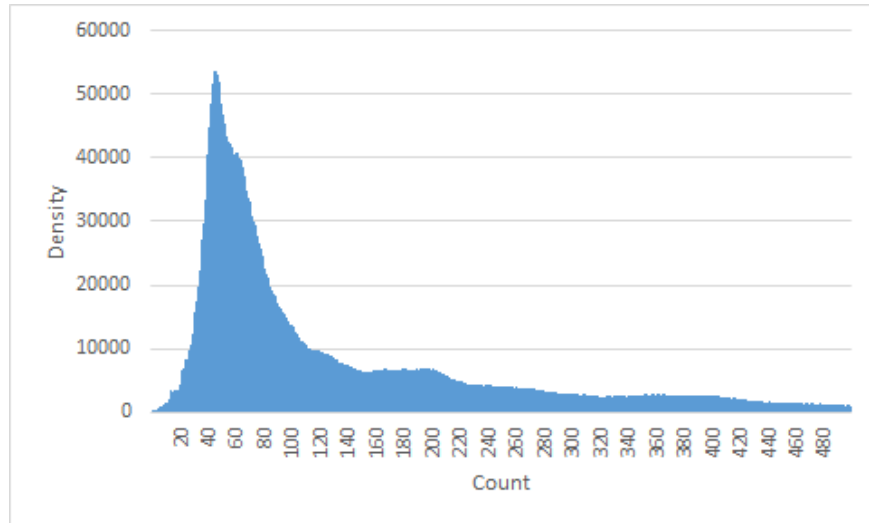


Figure 3 - Point Distribution per 0.5m Node

A significant portion of the grid nodes contains upwards of tens of thousands of soundings. Assuming the target node spacing of the survey was only 0.5m, this area was grossly over-sampled for the intended usage. Conversely, without a constraint on the target node spacing, a resolution much finer than 0.5m is achievable with this dataset. Returning to the fixed-resolution case, re-gridding the data at a 20cm grid spacing yields the following:

Method	Resolution	Creation Time	Contours	Sounding Selection	Total Time
- CUBE algorithm	0.2m	00:12:21	00:02:02	00:01:15	00:15:38

Table 3 - Scenario 1, Additional 20cm Gridding

At this point, the fixed-resolution surface at a 0.2m resolution takes approximately one quarter of the time to produce versus an unconstrained variable resolution surface over the same area, even though the 0.2m surface takes twice as long as the 0.5m.

Range Estimation

The Ranges estimation method, which simply assigns a resolution to an area based on a lookup table of depth range and resolution, was only slightly less performant than the fixed-resolution case. This allows the flexibility of Variable Resolution surfaces by applying some a-priori knowledge or expectations about the density of the source data. However, these a-priori assumptions result in the same problems the fixed-resolution surfaces, namely that our gridding resolutions of 0.5m and 1m grossly under-sample relative to the actual survey data density. Depending on the intended product, again we are either wasting costly acquisition time or not making full use of the data density in our products.



Figure 4 - Variable Resolution Colour by Resolution, 0.5m and 1.0m

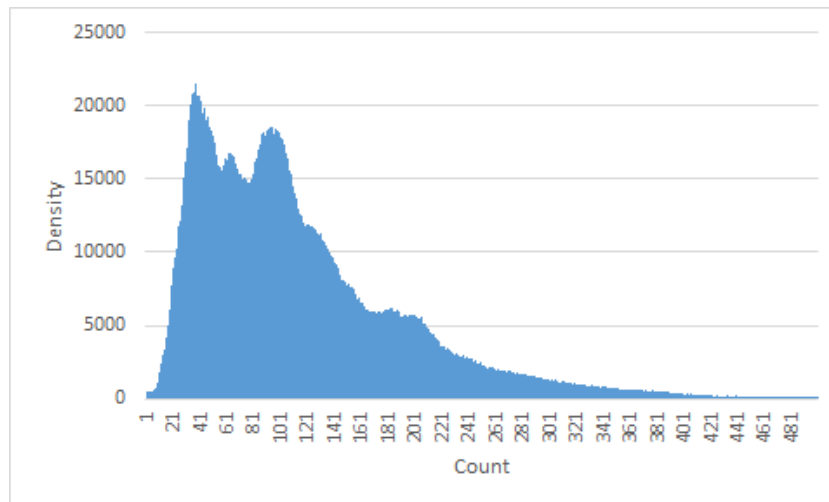


Figure 5 - Range Estimation Method, Point Distribution per Node

Discussion

Over a small, shallow area such as Plymouth Harbour, generating a Variable Resolution product using a density-based method may not be cost-efficient in terms of processing time compared to traditional fixed-resolution surfaces. If we have some a-priori knowledge of the intended product resolution, and are cognizant of this during acquisition in terms of our data density, the Ranges method will provide a preferred approach when dealing with surface vessel platforms.

Alternatively, value can be derived from a density analysis as well. It provides insight into our survey practices, specifically whether we are under- or over-sampling in a given survey area, relative to our expected data density target for a given survey product.

Assuming a fixed resolution is not a requirement for the final product, density estimation also offers the benefit of a best-fit resolution in discrete areas across our dataset, and compromising on resolution is not necessary. An example of this would be developing a

specific feature in a survey area: assuming a coarse resolution across the majority of a survey area, and fine detail over a feature of interest, the coarse and fine data could be represented as a single product when applying density estimation with Variable Resolution. With the added processing time, a realistic scenario may be generating a fixed-resolution grid for data QC and cleaning, and producing a Variable Resolution surface only once as a final product.

Following on to the scenario of developed features in a survey, in the case of the Shallow Survey 2015 dataset additional feature detection lines were run across specific targets of interest, in addition to the main-scheme survey lines. Running a new variable resolution surface with this extra data included, and plotting the resulting resolution, in Figure 7 we can see a finer result over areas with the extra feature detection lines compared to the original variable resolution surface in Figure 6.

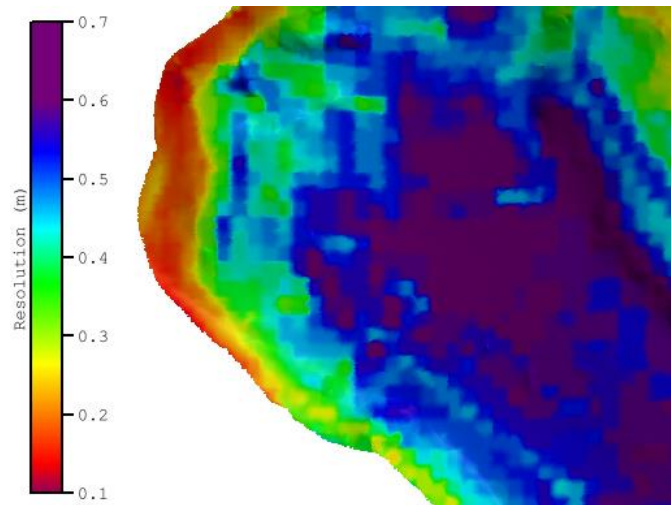


Figure 6 - Original VR Surface Coloured by Resolution




Figure 7 - VR Surface Coloured by Resolution, Feature Detection Lines Overlaid

Scenario 2

The dataset of interest for Scenario 2 is a NOAA survey conducted east of Glacier Bay National Park and Preserve, sheet H12142. This dataset was produced from multiple platforms (1 ship + 4 launches) over a reasonably broad depth range, <0m to approximately 420m. This combination of platforms and working depths is an ideal real world scenario to consider for application of Variable Resolution. Grids were generated using the standard resolution and depth ranges for NOAA surveys. Contours were generated at IHO-standard intervals of 0, 2, 5, 10, 20, 50, 100, 200, 300, 400m where these intervals overlapped the depth range. Sounding Selection was applied using a radius table where spacing increases with depth, against appropriate depth intervals.

- Sonars:
 - Ship: Reson 7111, 8160
 - Launches (4): Reson 7125, 8101
- 535,745,793 soundings
- 1,099 lines
- 147.2 km², 1,040 LKM
- Minimum depth: -3.0m
- Maximum depth: 426m




Figure 8 - H12142

Results

The broad depth range, up to 426m from 0m, required that 6 different single resolution surfaces be created, according to NOAA Specifications and Deliverables (NOAA 2016). Note that the time it takes to set up and run each process (6 times each or batch process set up) is not counted in the creation times. Timing results of the single resolution surfaces are as follows:

Method	Resolution	Creation Time	Contours	Sounding Selection	Total Time
- CUBE algorithm	1.0m	00:42:11	00:08:03	00:03:48	00:54:02
- CUBE algorithm	2.0m	00:32:31	00:01:07	00:01:26	00:35:04
- CUBE algorithm	4.0m	00:13:53	00:00:20	00:00:22	00:14:35
- CUBE algorithm	8.0m	00:11:34	00:00:06	00:00:06	00:11:46
- CUBE algorithm	16.0m	00:07:01	00:00:02	00:00:03	00:07:06
- CUBE algorithm	32.0m	00:05:55	00:00:00	00:00:00	00:05:55
Total time		01:53:05	00:09:38	00:05:45	02:08:28

Table 4 - Scenario 2, Fixed Resolution Processing Times

Timing results for the Variable Resolution surfaces are as follows:

Method	Resolution	Creation Time	Contours	Sounding Selection	Total Time
Variable Resolution					
- CARIS Density	Min 0.09m	01:22:59	00:33:33	00:01:25	03:06:44
- CUBE algorithm	Max 40.0m	01:08:47			
Variable Resolution					
- CARIS Density	Min 1.0m	00:29:31	00:15:59	00:00:46	01:25:26
- CUBE algorithm	Max 54.0m	00:39:10			
Variable Resolution					
- Ranges estimation	Min 1.0m	00:09:19	00:42:05	00:00:51	01:33:14
- CUBE algorithm	Max 18.0m	00:40:59			

Table 5 - Scenario 2, Variable Resolution Processing Times

Discussion

The unconstrained density estimation for Variable Resolution is measurably longer in processing time versus the creation of multiple single-resolution products, approximately 3 hours vs 2 hours. However, an important consideration to be made in this case is NOAA's practice of limiting their gridded resolutions to 1.0m, sometimes 0.5m. Similar to the findings in scenario 1, from the density estimation we can see a minimum of 9cm in the output surface, which is significantly smaller than 1.0m in the fixed-resolution datasets. Therefore, a more appropriate comparison to make in this instance is with the second density estimation case, where the minimum resolution is limited to 1.0m to match. In this case, Variable Resolution is measurably faster to produce the same product (~1.5 hours instead of 2). A key discovery can again be made from these findings: either significant portions of the area are grossly over-sampled, and acquisition efforts in future surveys can be reduced, or the products

produced from these surveys are not appropriately modeling the level of detail available in the source data.

Another consideration to make in fixed-resolution versus Variable Resolution processing is the practical case of data management. Through long-standing practices, a series of fixed resolutions and associated depth ranges are applied to produce between 1 and 6 grids for each survey. During data processing, each of these grids must be maintained simultaneously: as new data is acquired each day, it must be processed as normal and then added, either manually or through batch commands, to each of these grids. This process is prone to errors, omissions, missed or corrupted datasets. Often re-processing on specific portions of the dataset must be performed, perhaps to apply final tides or post-processed positioning data, and in this instance each grid must be regenerated again to pick up these changes. In the case of a full regeneration, this process takes the same amount of time as the original grid creation. If this must occur several times over the course of processing a complete survey, this is a significant time investment, easily compounded by user error. Creating and maintaining a single surface product, in this instance, greatly reduces the chance for user errors and simplifies a multi-week, multi-platform survey considerably.

Another aspect of such a large survey is the creation of vector products. For the single resolution surfaces, contouring and sounding selection are run on each surface dataset. Where these datasets meet, there is currently no way to have consistent feature edge mapping for contours (especially) and soundings, leading to hand-drawn contours or manual editing of repeated soundings.




Figure 9 – Contours and Soundings from Overlapping Grids (Red and Blue)

Having a single Variable Resolution product from which to create feature products completely eliminates the edge mapping issue, leading to consistent features across the entire survey area.

Scenario 3

Scenario 3 consists of a combined multibeam sonar and bathymetric LiDAR survey at Shilshole, Seattle, WA. This dataset was collected by Fugro Pelagos and provided to CARIS for testing and analysis. It provides a unique combination of high-density multibeam with overlapping, lower-density bathymetric LiDAR across an in-shore area. All noted depths are referenced to the WGS 84 ellipsoid.

Multibeam dataset:

- Sonar: Reson 8101
- 21,923,767 soundings
- 84 lines
- 1.33 km², 50.6 LKM
- Minimum depth: 0.821 m
- Maximum depth: 55.405 m

LiDAR dataset:

- Scanner: Shoals 1000T
- 456,568 shots
- 12 lines
- 1.05 km², 31.6 LKM
- Minimum depth: -41.774 m
- Maximum depth: 19.131 m




Figure 10 - Shilshole Points, Multibeam on the Left, LiDAR Centre/Right

The multibeam dataset was collected in April of 2005, and the LiDAR was flown in August of 2007. Due to the age of the surveys, a complete uncertainty model was not provided with the data. For comparison purposes relative to the other use cases, a rational uncertainty model was assembled through some reasonable estimates based on sensors available at the time to allow computation of Total Propagated Uncertainty

(TPU) values, such that the CUBE gridding algorithm could be applied in processing. For the purposes of this analysis, this estimated uncertainty model was deemed a reasonable approach. Contours were generated at intervals 0, 10, 20, 30, and 40m (-20, -10, 0, 10 and 20m for the LiDAR and combined datasets), and the Sounding Selection method was done using a radial spacing of 5.0m on the ground.

Results

We will quantify processing the multibeam and LiDAR datasets as separate products, and also combined. Fixed resolution timing results are as follows:

Method	Resolution	Creation Time	Contours	Sounding Selection	Total Time
- CUBE algorithm - Multibeam data	1.0m	00:01:37	00:00:18	00:00:20	00:02:15
- CUBE algorithm - LiDAR data	10.0m	00:00:04	00:00:05	00:00:02	00:00:11

Table 6 - Scenario 3, Fixed Resolution Timings

And Variable Resolution timings:

Method	Resolution	Creation Time	Contours	Sounding Selection	Total Time
Variable Resolution - Multibeam only			00:01:06	00:00:14	00:09:37
- CARIS Density	Min 0.38m	00:05:01			
- CUBE algorithm	Max 18.24m	00:03:16			
Variable Resolution - Multibeam only			00:00:54	00:00:15	00:03:50
- Ranges estimation	Min 1.0m	00:01:03			
- CUBE algorithm	Max 4.0m	00:01:38			
Variable Resolution - LiDAR only			00:00:04	00:00:04	00:00:24
- CARIS Density	Min 8.0m	00:00:09			
- CUBE algorithm	Max 14.7m	00:00:07			
Variable Resolution - Combined dataset			00:01:27	00:00:24	00:12:17
- CARIS Density	Min 0.4m	00:06:51			
- CUBE algorithm	Max 27.3m	00:03:35			

Table 7 - Scenario 3, Variable Resolution Timings

The multibeam sonar and bathymetric LiDAR datasets were each gridded separately at 2 different single resolutions (1m and 10m, respectively). Figure 11 shows these surfaces grouped together (for display purposes).




Figure 11 - Grouped Fixed Resolution Display

The Variable Resolution surface produced shows a single gridded product that encompasses both datasets:




Figure 12 - Combined Variable Resolution Surface

We can also colour the resulting surface by resolution:




Figure 13 - Variable Resolution Surface, Coloured by Resolution

It is here we begin to notice some oddities over some of the LiDAR data regions. Inspecting the lower-left corner of the surface at a larger scale begins to show the problem:



Figure 14 - Variable Resolution Surface Artifacts

As with the 2 single resolution surfaces from each dataset, shown above in Figure 11, we would expect the gridded data to cover the entire survey area. Using the Variable Resolution with CARIS Density estimation, large gaps in the gridded areas are apparent.

The problem becomes somewhat clearer by overlaying the resolution map, which shows the layout of tiled regions, each assigned a single fixed resolution.




Figure 15 - Variable Resolution Surface with Resolution Map Overlay

The blue lines indicate tile boundaries, and black lines are the individual grid node boundaries within each tile. To the north we have the dense multibeam, interspersed with the coarser LiDAR data across the south. The region shaded in blue is at a much finer node spacing than the surrounding tiles. The sounding data within the (shaded blue) tiles with the finer node spacing does not support this resolution over the entire tile, leading to gaps in coverage (artifacts).

Discussion

As also seen in scenario 1, over a small area the creation of an initial Variable Resolution surface (CARIS Density method) can take much longer in post-processing than the fixed-resolution case. With some a-priori knowledge of the data density and/or required product resolution, the density estimation could be constrained more aggressively, or the Ranges method used, to save on processing time. However, it is more relevant to apply Variable Resolution to these datasets because of the sharp disparity in data density between the multibeam and LiDAR data; the multibeam data is quite dense, easily supporting below a 1m resolution in shallow areas, whereas the LiDAR data on average meets a 10m spacing. For the fixed resolution scenario, we must either treat these as two neighboring, yet independent, products, or grid the entire survey at the coarser 10m resolution, losing much fine detail where the multibeam has coverage. With Variable Resolution we endeavor to evaluate both surveys at an appropriate resolution in a single product.

However, upon inspection of the pattern of artifacts in the surface produced by the Variable Resolution CARIS Density method (Figure 15), it becomes apparent that the tile structure of Variable Resolution can be limiting when we have regions of fine data density bounding other regions of coarse data density. This abrupt density change can overwhelm the density estimation, which is only ever evaluating a squared region, and the estimation will generally favour the finer resolution (Beduhn, Foster, MacGillivray 2016). This results in artifacts in tiles containing both data sources, when the density

estimation determines a fine resolution is appropriate for a tile while also gridding the coarse data in the same region at that fine resolution.

A more optimal approach in this case may be to allow these generalized regions of varying density to be defined a-priori, and feeding that into the density estimation algorithm such that the two datasets are treated independently while still allowing a single continuous dataset stored as a result. The Variable Resolution architecture allows for such a structure to be stored, but at this time the interactive tools and processes are not in place to allow it. Work is now under way to support this more challenging data distribution.

Specifications and Deliverables

Surveys such as these, particularly scenario 2, are ultimately used for nautical chart updates, therefore adequate feature detection and data coverage are both equally important.

The International Hydrographic Organization S-44 standard (2008) states, "when a full sea floor search is required, the equipment used to conduct the survey must be demonstrably capable of detecting features of the dimensions specified." Features are navigationally significant objects, and for example, S-44 Order 1a, defining areas "where the sea is sufficiently shallow to allow natural or man-made features on the seabed to be a concern," specifies a minimum 2m cube object should be detectable by a survey system. "It is the responsibility of the hydrographic office / organization that is gathering the data to assess the capability of any proposed system and so satisfy themselves that it is able to detect a sufficiently high proportion of any such features" (IHO 2008).

The question is, then, are Variable Resolution surfaces adequate for feature detection? In all three scenarios above, both resolution estimation methods for Variable Resolution, CARIS Density and Depth Range, were capable of creating very high resolution models (e.g. <1m) from sources which support that resolution. In the fixed-resolution scenario, we can generate a 1m grid and ensure we have an adequate number of samples per node, with no gaps in that surface. However, in the Variable Resolution scenario, our resolution spacing is no longer fixed. Per the NOAA NOS Hydrographic Surveys Specifications and Deliverables (2016), "at least 95% of all nodes on the surface shall be populated, with at least 5 soundings," in order to meet object detection coverage. An important consideration in the creation of Variable Resolution surfaces using Density estimation is we define a target point density, which is a general target density for all nodes in a given tile. To ensure we produce a minimum count (e.g. 5 per node), we can simply perform a logical query of the surface attributes to locate regions that do not meet this requirement. At the same time, we can also query the resolution of each node also meets the appropriate target.

The following calculation ensures each node has a minimum of 5 samples, and that the resolution is at least 1m up to 40m depth, and 10% of water depth beyond that:

```
IF (Density >= 5 AND ((Depth<=40 AND Resolution<=1) OR (Depth>40 AND Resolution <= (0.1 * Depth)))) THEN NO_DATA ELSE 1
```

Using the Compute Layer function on the Variable Resolution surface, we can quickly determine if we have areas in the dataset that do not meet this criteria:




Figure 16 - Selection of Nodes that do not meet IHO Criteria

From this selection, we know we have 10,084 nodes failing of 10,629,753 total nodes or ~0.1%, well within the maximum 5% of failing nodes.

The IHO S-44 Specification (2008) also has a concept of "full seafloor coverage," which is a systematic method of exploring the seafloor in order to detect all features, as defined by the S-44 Orders. NOAA (2016) takes this a bit further by defining what a "holiday" (gap) is with respect to their defined object detection coverage. In this way they can systematically look for holes in the grid meeting this criteria and make decisions about how to treat them (e.g. resurvey for coverage if critical underkeel clearance is needed or ignore if not). This method is very straight forward for single resolution surfaces, as the resolution chosen indicates a depth range and subsequently the IHO Order requirement for that area of the survey.

Therefore, the second question is, can Variable Resolution surfaces adequately demonstrate the concept of "full seafloor coverage" (IHO 2008), or in NOAA's case, object detection coverage with no holidays? Locating data holidays, i.e. gaps in the surface, becomes an interesting problem with Variable Resolution surfaces. Consider the fixed-resolution case:




Figure 17 - Examples of Detection Holidays, in red, NOAA (2016)

Holidays are defined in terms of their grid resolution. So at a given grid spacing (1 m , 4 m , 8 m etc.), any 3 collinear empty nodes are considered holidays. Association a 3-count with a resolution, we have gap size of 3 m , 12 m , 24 m etc. We can reasonably assume this is the maximum holiday requirement within the depth range associated with those resolutions. Therefore, the process seems straight-forward for Variable Resolution that we look for holes over a certain size, rather than a number of nodes. However, consider the following example between two tiles of differing resolution:




Figure 18 - Variable Resolution Surface Holidays

In the above two scenarios, assuming less than 40 m water depth, are they both Holidays (3 m gap)? How is that gap measured? A $3 \times 3\text{ m}$ cube would likely be detected in this case, however a linear feature 3 m in length may not. It is still uncertain at this time what the solution may be for determining holidays under a variable resolution paradigm.

Conclusion

Variable Resolution can make the creation of gridded products across regions of varying source density much easier, by creating a single product instead of many, as seen in scenario 2. Both scenario 1 and 2 highlight that hydrographic surveys are in many cases collecting data at densities much finer than their deliverables require. Data can therefore be acquired at reduced settings and likely higher acquisition rates, or much higher-resolution products can be realized from these surveys than is currently done in practice. Scenario 3 highlights that abrupt changes in data density, such as this case between multibeam and bathymetric LiDAR, can be problematic for density estimation algorithms that consider only squared regions as a whole, and have no consideration for these abrupt and non-linear density changes. A gap still exists under this scenario, but a few adaptations of the existing toolset will likely serve to solve this problem as well, as the fundamental data structures do support storage of these arbitrary regions. Finally, many existing specifications can be adapted such that variable resolution surfaces can meet their requirements, although some challenges still exist.

The algorithms used to create the resolution maps for Variable Resolution surfaces have been thoroughly vetted, and progress is being made on further improving them and even creating new ways to automatically determine appropriate resolutions from sounding data. Considerations have also been made for further downstream processing and data management. Other CARIS products can also generate VR surfaces, and they can be stored in a bathymetric database. Products can be generated directly from these surfaces whether in HIPS, BASE Editor or stored in the database. This technology is scalable and it will play a major role in advancing the capacity of hydrographic survey and general elevation data analysis, by providing an efficient way to process and manage big data.

References

Beduhn, T., Foster, B., MacGillivray, S., 2016, Applications of Resolution Estimation in Variable Resolution Surfaces, Canadian Hydro 2016, Halifax, NS, Canada, 16-19 May. Conference Proceeding

International Hydrographic Organization (IHO). 2008. IHO Standards for Hydrographic Surveyors, Special Publication No 44. 5th Ed.: 36pp.

National Oceanic and Atmospheric Administration (NOAA). 2016. NOS Hydrographic Surveys Specifications and Deliverables. 194pp.

List of Figures

Figure 1 - Plymouth, UK, SS2015	5
Figure 2 - Resolution Distribution in a Variable Resolution Surface	6
Figure 3 - Point Distribution per 0.5m Node	7
Figure 4 - Variable Resolution Colour by Resolution, 0.5m and 1.0m	8
Figure 5 - Range Estimation Method, Point Distribution per Node	8
Figure 6 - Original VR Surface Coloured by Resolution	9
Figure 7 - VR Surface Coloured by Resolution, Feature Detection Lines Overlaid	9
Figure 8 - H12142	10
Figure 9 – Contours and Soundings from Overlapping Grids (Red and Blue)	12
Figure 10 - Shilshole Points, Multibeam on the Left, LiDAR Centre/Right	13
Figure 11 - Grouped Fixed Resolution Display	15
Figure 12 - Combined Variable Resolution Surface	15
Figure 13 - Variable Resolution Surface, Coloured by Resolution	16
Figure 14 - Variable Resolution Surface Artifacts	16
Figure 15 - Variable Resolution Surface with Resolution Map Overlay	17
Figure 16 - Selection of Nodes that do not meet IHO Criteria	19
Figure 17 - Examples of Detection Holidays, in red, NOAA (2016)	20
Figure 18 - Variable Resolution Surface Holidays	20

List of Tables

Table 1 - Scenario 1, Surface Timings	5
Table 2 - Scenario 1, Additional Density Analysis	6
Table 3 - Scenario 1, Additional 20cm Gridding	7
Table 4 - Scenario 2, Fixed Resolution Processing Times	11
Table 5 - Scenario 2, Variable Resolution Processing Times	11
Table 6 - Scenario 3, Fixed Resolution Timings	14
Table 7 - Scenario 3, Variable Resolution Timings	14

Trademarks owned by Teledyne CARIS, Inc.

This is a listing of USPTO-registered trademarks and trademarks owned by Teledyne CARIS, Inc. and might also be trademarks or registered trademarks in other countries. Please note that laws concerning use and marking of trademarks or product names vary by country. Consult a local attorney for additional guidance. Teledyne CARIS, Inc. permits the use of its trademarks and registered trademarks only where they are used in reference to Teledyne CARIS, Inc. and its products, the markings used are appropriate to the country or countries of publication, and Teledyne CARIS, Inc. is explicitly acknowledged as the owner of the mark. Teledyne CARIS, Inc. reserves the right to withdraw this permission at its sole discretion for any use it feels is inappropriate or adverse to its interests. Teledyne CARIS, Inc. otherwise prohibits the use of any of its registered symbols, insignia, or other identifying marks without express written approval. Violations are subject to liability for damages, injunctive relief, attorney's fees and other penalties.

Not all trademarks used by Teledyne CARIS, Inc. are listed in this document. Failure of a mark to appear on this page does not mean that Teledyne CARIS, Inc. does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market. The absence of a product or service name or logo from this list or the absence of a TM or TM Reg. USPTO notation against a product or phrase listed below does not constitute a waiver by Teledyne CARIS, Inc. of its trademark or other intellectual property rights concerning that name or logo.

The following are trademarks or USPTO-registered trademarks of Teledyne CARIS, Inc.:

- Article 76 Module
- Bathy DataBASE
- Bathy DataBASE Server
- BASE Editor
- BASE Manager
- BDB
- CARIS
- CARIS GIS
- CARIS Notebook
- CARIS Onboard
- ChartServer
- CPD
- Core Production Database
- Easy View
- EAM
- Engineering Analysis Module
- HIPS
- HIPS and SIPS
- HPD
- HPD Server
- Hydrographic Production Database
- Limits and Boundaries Module
- LIN
- LOTS
- LOTS Browser
- LOTS Limits and Boundaries
- LOTS Article 76
- One Feature, One Time
- Paper Chart Composer
- Paper Chart Editor
- Ping-to-Chart
- Product Editor
- Publications Module
- S-57 Composer
- SIPS
- Source Editor
- Spatial Fusion
- Spatial Fusion Enterprise
- Teledyne CARIS

Those trademarks followed by or footnoted as ^{TM Reg. USPTO} later in this document are registered trademarks of Teledyne CARIS, Inc. in the United States; those followed by or footnoted as ^{TM Reg. CIPO} are registered trademarks of Teledyne CARIS, Inc. in Canada; those followed by or footnoted as either ^{TM Reg. USPTO and CIPO} or ^{TM Reg. USPTO, CIPO} are registered trademarks of Teledyne CARIS, Inc. in both the United States and Canada; those followed by or footnoted as TM are trademarks or common law marks of Teledyne CARIS, Inc. in Canada and the United States, and in other countries.

The trademarks and names of other companies and products mentioned herein are the property of their respective owners.

Copyright owned by Teledyne CARIS, Inc.

All written and image content in this document not protected by the copyrights of others is © Copyright 1992 Teledyne CARIS, Inc. All rights reserved. All reproduction and redistribution is strictly prohibited without the express prior written consent of Teledyne CARIS, Inc. Copyright © 3/6/17 Teledyne CARIS, Inc. All rights reserved.