

Recent Improvements to Seafloor Imagery Acquisition and Processing Procedures for R2Sonic Multibeam Echosounders

J. Beaudoin¹, T. Reis² and M. Doucet³

1: Jonathan Beaudoin, Affiliate Assistant Professor
Center for Coastal and Ocean Mapping
University of New Hampshire
24 Colovos Road, Durham, NH, USA
jbeaudoin@ccom.unh.edu

2: Tom Reis, President
Substructure, Inc., P.O. Box 4094, Portsmouth, NH, USA
tom@substructure.com

3: Maurice Doucet, Chief of Software Research
QPS Inc., A Saab Defence and Security Company
104 Congress St., Suite 304
Portsmouth, NH, USA
doucet@qps-us.com

Abstract

Bathymetric performance is typically the primary focus when bringing a new multibeam echosounder to market. Whereas software vendors and seabed mappers quickly adjust to new hardware products, support for a system's seabed imaging capabilities often lags behind. Introduced to the market in 2008, the R2Sonic 2024 multibeam echosounder is no exception. The 2024 quickly established itself as a hydrographic-grade mapping system, however, limited understanding and support in post-processing software, along with a general lack of configuration and acquisition "best practices" knowledge in the mapping community, has limited its widespread use for projects where seabed imagery products are required.

In this paper, we examine the technical challenges involved with establishing an acquisition and processing workflow for R2Sonic 2024 multibeam echosounders for a large-scale mapping project of Lake George, NY. For this project, two vessels operated for nearly two months in late 2013 with one of the vessels being outfitted with a dual head configuration.

Challenges included:

- Implementation of a real-time method for monitoring signal saturation
- Determining the ideal dual-head frequency separation to avoid interference
- Establishment of a backscatter reference surface to support inter-vessel calibration
- Establishment of acquisition guidelines and best-practices
- Automation and streamlining of CTD data delivery for refraction and attenuation corrections

- Adding support for dual-head post-processing
- Streamlining data import procedures
- Improvement of imagery quality in post-processing

A close partnership between surveyors, academic consultants and hardware/software vendors allowed not only for the successful completion of the project but has also demonstrated an acquisition and processing workflow where none existed before.

Introduction

The Jefferson Project at Lake George is a three-year collaborative research effort between Rensselaer Polytechnic Institute, The Fund for Lake George, and IBM Research. Substructure Ltd of Portsmouth, NH, USA, was awarded a contract by IBM in 2013 to map bathymetry of Lake George, NY, USA. Lake George measures 114 km² and has a mean depth of 21 m and maximum depth of 76 m. Work was to proceed over a 2 month period in the fall of 2013, with two vessels operating daily: MV *Orion* with a single headed multibeam system to map the deeper portions of the lake and MV *Mintaka*, equipped with a dual head multibeam system to gain coverage in the shallow areas along the shoreline. Operations were suspended after the winter freeze up in early January 2014 and were resumed in Spring 2014.

Identical survey system configurations were made on both vessels to ease training, operation and post-processing for the survey crew. System components are listed below.

- Hardware:
 - R2Sonic 2024 multibeam echosounders, 1.0° transmitter, 0.5° receiver, 256 beams, 200-400 kHz
 - Applanix POSMV: RTK aided, SBET processing
 - YSI CastAway CTD
- Software:
 - QPS QINSy for acquisition, bathymetric cleaning/processing
 - QPS FMGT for backscatter processing
 - QPS Fledermaus for visualization

Bathymetric data were the main deliverable, however, high quality backscatter imagery was highly desirable as well. Though Substructure has successfully used their preferred survey configuration for bathymetry on multiple projects prior to the Lake George mapping effort, there were many technical challenges to be overcome regarding the backscatter. In large part, the challenges come from the fact that there is little community knowledge regarding how best to configure, acquire and process seabed imagery from R2Sonic echosounders.

Perhaps the largest hurdle in improving community knowledge has been a lack of commercially available software that could adequately process R2Sonic imagery data. Without software to adequately process data, surveyors and mappers have not been able to experiment with survey configurations and develop appropriate best practices for seabed imagery acquisition and processing for R2Sonic systems. Previous unpublished work done in 2011 by Beaudoin led to the development of R2Sonic processing algorithms in the University of New Brunswick's (UNB) SwathEd software; these algorithms were provided to Caris and QPS in 2011 and 2012, but problems with these implementations still remained and the community in general did not have a viable post-processing workflow. To address this issue, a partnership was forged:

- **Substructure, Ltd:** Surveyor willing to take on risk and experiment with new tools and methods.
- **University of New Hampshire:** Researchers willing to assist in knowledge transfer, survey and best-practice design and in industrialization of research code.
- **QPS:** Software vendor willing to implement research code into commercially available solution.
- **R2Sonic:** Hardware vendor willing to provide information and support.

It is the intent of this paper to discuss, at a broad and general level, the technical challenges encountered and overcome in this work. The following topics are addressed:

1. Monitoring RX Signal Saturation
2. Optimizing Dual-Head Frequency
3. Backscatter Reference Surface
4. Streamlining CTD input
5. Acquisition Guidelines & Best-Practices
6. Improving Post-Processing

With (1) and (2), it was necessary to establish whether or not these could even be done prior to mapping efforts. With (3), (4) and (5), procedures were implemented to aid in the ability to consistently acquire high quality data that minimized the level of effort in post-processing. With the field procedures in place, the last hurdle was to ensure that the quality of imagery products that were possible with UNB SwathEd could be replicated in commercial software, namely QPS FMGT.

Challenges

Monitoring RX Signal Saturation

One of the primary limiting factors in terms of seabed imagery fidelity for some multibeam echosounders is the limited dynamic range of the analog-to-digital (A/D) convertors. With limited dynamic range, the analog signal needs to be adjusted prior to digitization, typically with a time-varying gain (TVG) to ensure that the analog signal does not (1) fall below the noise floor of the A/D or (2) extend above the A/D's maximum input voltage limit.




Figure 1. Example of signal saturation using a WASSP multibeam echosounder system in Portsmouth Harbor, NH, USA. Multiple passes over the same seabed were run with differing power levels to demonstrate the onset of saturation. Note the change in signal response in the angular response curves in the lower image, in these cases, the saturation has occurred at nadir for the passes with higher power levels whereas the data from the outermost beams are acceptable for all passes. Image courtesy of Dr. Tom Weber (UNH-CCOM).

The first problem is less of a concern since the bottom would not be detected in this particular scenario and the surveyor will react immediately to rectify it, e.g. increase the power or pulse width. With the second problem, the sonar will still detect the bottom and imagery will continue to be output, however the receiver signal has saturated. Saturation distorts the true signal level beyond recovery and the system is no longer reporting the true echo level, i.e. the echo level will

not increase/decrease in response to changes in seafloor reflectivity. Furthermore, when the system is operated in a manner which results in saturation, corrections for power and gain are meaningless since the output signal level no longer increases accordingly and post-processing procedures introduce artifacts where there were none before (though the underlying signal is still saturated). Saturation can occur in any operational configuration or mapping environment where the echo level becomes strong enough to exceed the A/D dynamic range, for example, with high power levels or excessive gain in shallow water with highly reflective seabed (see Fig. 1). The reader is referred to Rice et al. (2010) for more discussion on the causes and effects of saturation.

Since the true signal level cannot be recovered in post-processing once saturation has occurred, there is a clear need to monitor for saturation and to react accordingly in the field. Initial fieldwork done by Beaudoin and Reis confirmed that the R2Sonic systems are susceptible to saturation and that a real-time saturation monitor was required for the Lake George mapping project.

Based on research done by S. Greenaway at UNH, Rice had developed Saturation Monitoring tools for NOAA's Reson multibeam systems (Rice et al., 2010). This implementation of a saturation monitor was used as a model for the R2Sonic monitor, both in its approach to monitoring saturation in real-time and also in construction of the saturation curve, which is dependent on the receiver gain and TVG. Beaudoin consulted with Rice and implemented an independent and automated solution to determine R2Sonic saturation curve characteristics for R2Sonic systems. Beaudoin then worked with R2Sonic developers to create a real-time monitor that used the saturation curve characteristics measured by Beaudoin and the beam specific intensity and TVG values to determine saturation susceptibility. After some field-testing in Tampa Bay, FL, USA with C. Brennan of R2Sonic, a new version of the R2Sonic Sonic Control module was delivered to Substructure in time for the Lake George mapping project. Beaudoin joined the Substructure field crew to install and test the software in addition to providing training to the crew. Figure 2 shows a photograph of the R2Sonic saturation monitor as installed on MV *Mintaka*, the dual-headed vessel.




Figure 2. Photograph of the R2Sonic saturation monitor as installed on MV *Mintaka*, the vessel with the dual-headed configuration. Saturation is detected when the beam signal levels approach or exceed the red line signifying the saturation limit. The operator can avoid saturation in a number of ways: (1) reducing the power, (2) reducing the gain, preferably by adjusting the TVG spreading parameter.

Optimize Dual-Head Frequency

The second problem that involved optimization of operational parameters was that of determining the optimal frequency configuration for the dual head system. Seeing as seafloor backscatter is frequency dependent, it was desirable to minimize the frequency separation between the two vessels, MV *Orion* and MV *Mintaka*, in addition to minimizing the frequency separation between the two sonar heads on MV *Mintaka*. This is done to minimize differences due to the seafloor response to different frequencies in the overall mosaics.

The simplest method to achieve this is to use the same frequency for both heads on MV *Mintaka* and to configure the heads for interleaved pinging, sometimes referred to as a ping-pong configuration, where one head pings and completes its reception cycle, followed by the second head pinging and completing its reception cycle. This configuration was not an option due to the constraint it would impose on survey speed since the lower ping rate would halve the along-track sounding spacing and the vessel speed would need to be reduced significantly to overcome this.

Simultaneous pinging was required, thus the two sonar heads installed on Mintaka had to ping with a slightly different frequency in order to avoid interfering with one another. It was also desired to use high bandwidth pulses (short pulses) for the best resolution possible. High bandwidth pulses require more frequency separation between the sonar heads in order to avoid interference. The desire for high-resolution was balanced with the need to minimize the frequency separation by finding the absolute minimum separation in frequency possible for a desired pulse width. The minimum required separation was calculated based on pulse bandwidth and was tested in the field by Beaudoin and C. Brennan (R2Sonic) in Tampa Bay, FL, USA prior to the project. The recommended frequency separation calculation for central frequency f_c :

$$\Delta f = 1.4 \times BW_\tau = 1.4/\tau$$

$$f_1 = f_c + 0.5\Delta f$$

$$f_2 = f_c - 0.5\Delta f$$

These equations were used to prepare a look up table for the Substructure mapping crew to assist in selection of appropriate frequencies based on the pulse length used in operation (see Table 1).

Table 1. Frequency Separation for Dual Head Configuration

Pulse length (μs)	f₁ (kHz)	f₂ (kHz)
25	328.0	272.0
30	323.3	276.7
35	320.0	280.0
40	317.5	282.5
45	315.6	284.4
50	314.0	286.0
55	312.7	287.3
60	311.7	288.3
65	310.8	289.2
70	310.0	290.0
75	309.3	290.7
80	308.8	291.3
85	308.2	291.8
90	307.8	292.2
95	307.4	292.6
100	307.0	293.0

Backscatter Reference Surface

*“A man with one watch knows exactly what time it is.
A man with two watches is never sure.”*

Due to limitations in available software to post-process R2Sonic imagery data, the repeatability and stability of their imagery measurements is not well understood in the mapping community. There have been known issues with this particular type of problem with other multibeam hardware vendors and it was important to either verify the repeatability of R2Sonic systems in general or at least to control for the potential for day-to-day variation in the output signal level of the three R2Sonic 2024s that were deployed on the two vessels.

To this end, a backscatter reference surface was established near the dock where both vessels were moored every day (see Fig. 3). This was done during the initial testing of the Saturation Monitor, thus system settings were found such that the line could be run comfortably below the saturation point of the systems. This gave a simple, yet effective, means of addressing system repeatability and stability over time and also addressed inter-system comparability. A short survey line was planned over the surface and the mapping crews were instructed to image this line every day before field operations with the following instructions:

- Acquire a CTD cast prior to running the lines, calculate and upload the attenuation coefficient from the CTD to both vessels.
- Always run the line from north to south.
- Use common sonar settings for both vessels:
 - Range: 40 m
 - Frequency:
 - 300 kHz (MV *Orion*)
 - 272 kHz & 328 kHz (MV *Mintaka*)
 - Power: 200 dB
 - Gain: 1 dB
 - TVG Spreading: 20 dB
 - Pulse length: 25 μ s
 - Angular sector: 120°

The data acquired for this have not been analyzed but the overall mosaic results indicate that the systems provided stable and repeatable measurements over the field campaign and that there was a fixed offset signal level offset of approximately 4 dB between MV *Orion* (single head) and MV *Mintaka* (dual-head).




Figure 3. Backscatter reference surface established near the vessel mooring location.

Streamlining CTD input

The Castaway CTD, manufactured by YSI, was used to measure conductivity, temperature and depth (CTD) profiles for both vessels with one CTD unit per vessel. The CTD data provided temperature and salinity estimates for computing the speed of sound, a necessary parameter for accounting for the effects of acoustic refraction on multibeam echosounder data. The CTD measurements also provided the necessary information to compute the absorption coefficient for the lake water, a parameter that allows for estimation of the transmission losses that the acoustic signal undergoes while traveling from the transmitter to the seafloor and back.

Given that the project was to extend over the two month period prior to freeze up, it was expected that the lake temperature would vary significantly over the mapping project and these effects would slowly but surely introduce bias into the backscatter measurements over the duration of the project unless these effects were accounted for in real-time or in post-processing. The absorption coefficient is one of the parameters that is typically used to configure the real-

time TVG in most multibeam echosounders and the value selected in real-time is nearly always stored in the multibeam data files, thus there was a strong case for pursuing a strategy which allowed for the input of this parameter in real-time.

The Multibeam Advisory Committee's (MAC) SVP Editor program (MAC, 2013) was chosen to bridge the gap between the CTD download software, the multibeam control software and the multibeam acquisition software, Sonic Control and QINSy, respectively. The driving need was to reduce the complexity of the backscatter data processing workflow since the CTD information would not need to be imported into backscatter post-processing if the information was correctly applied to the imagery data in the first place. The MAC SVP Editor already provided the capability to import YSI Castaway CTD files and to broadcast the resulting sound speed profile via UDP transmission to the multibeam acquisition software (QPS QINSy), however, slight modifications were required to compute the frequency dependent absorption coefficient and to also broadcast this information over the Ethernet to the R2Sonic control software, Sonic Control, which also required modifications to enable it to receive this information via UDP broadcast.

With these software modifications in place, a workflow was then provided to the Substructure field crew (see Fig. 4):

1. CastAway transmits to CTD processing software via BlueTooth
2. CastAway data are exported in .cnv format
3. SVP Editor opens the .cnv file
4. User verifies and edits as necessary, e.g. extending the cast, removing outliers
5. User transmits the SVP and CTD information to QINSy and Sonic Control where it is applied immediately without further human interaction

This workflow reduced human error during file upload/download and also sped up field operations since the entire process from acquisition to application took just a few minutes at most. Because the absorption coefficient was provided to Sonic Control, it was immediately used to update the TVG parameters for the sonar and was thus embedded into the R2Sonic data files such that the absorption coefficient used in real-time was available in post-processing and there was no need to import the CTD files into the backscatter processing software, QPS FMGT.




Figure 4. Schematic diagram of CTD workflow from acquisition to application of the data in Sonic Control (absorption coefficient) and in QINSy (for refraction corrections).

Acquisition Guidelines & Best-Practices

The previous sections introduced several best-practices for the Lake George project that are worth summarizing:

1. Use the saturation monitor to detect saturation and react accordingly.
2. The pulse length should be 25 μ s for both vessels during survey; set the dual-head operating frequencies to 272kHz and 328 kHz and the single-head to 300 kHz and do not adjust these.
3. Image the Backscatter Reference Surface daily with both vessels with exact same settings.
4. For every CTD, transmit the absorption coefficient to Sonic Control. Do not adjust the absorption coefficient in Sonic Control afterward for any reason.

A few aspects of sonar configuration and control remain which must be addressed, in particular since multibeam systems like the R2Sonic provide the operator with very granular control over sonar signal parameters such as power, gain, pulse length and frequency. To achieve excellent backscatter imagery results, the common practice with these types of systems is to “leave the

knobs alone” since the ability to correct the output signal level for the variations associated with these parameter changes is sometimes imperfect for a number of reasons. This practice, however, can be detrimental to the bathymetric data quality in terms of increased noise levels in the bottom detections and/or decreased swath width. In an area of complex topography such as Lake George where the bottom depth can reach up to 76 m, “leaving the knobs alone” is an imprudent course of action. Field testing was done with the R2Sonic 2024 to establish which, if any, of the various sonar parameters could be adjusted without introducing backscatter imagery artifacts. By design, the Substructure mapping crew had been instructed not to modify the pulse length, frequency and TVG absorption coefficient thus the only remaining parameters of interest that influence the output signal level are gain, TVG “spreading” and power. Through field testing in Tampa Bay, FL, USA, it was determined that corrections for real-time adjustments to these parameters can be adequately applied in post-processing but with the following caveats and/or additional best-practices:

- **Gain:** Can use full range of gain but be careful to avoid saturation.
- **TVG “Spreading”:** Adjustments to the TVG spreading parameter provide a good way to quickly avoid saturation when working against the constraints imposed with low power and power hysteresis artifacts (see below).
- **Power:** Two types of power related problems were found to introduce artifacts in backscatter imagery (discussed below) and the power should be adjusted with care due to avoid power hysteresis and low power sawtooth artifacts.

Power Related Artifacts: Hysteresis

Hysteresis is the phenomenon in which the value of a physical property, in this case the reported echo level, lags behind changes in the effect causing it, in this case the sonar power level. In the case of the R2Sonic multibeam echosounder, it was found that decreases in power level take several tens of pings to drop down to the requested level. On the other hand, increases in power are effective immediately and hysteresis was discernible in the case of adjusting the power upward. The lag in signal level reduction associated with power level decreases can introduce subtle artifacts in imagery when done carefully, however, dropping the source level several steps at a time introduces very obvious artifacts in imagery which persist for a much longer time and thus over a greater distance along the survey line (Fig. 5). At the time of the survey efforts, there was no mechanism to correct for this in post-processing, thus the best-practice recommendation to the field crew was avoid dropping the power level several steps at a time though it was perfectly acceptable to adjust the power upward without penalty.




Figure 5. Example of hysteresis type artifacts associated with reduction in power level in imagery data. The modest reduction in power introduces a subtle artifact on the left side whereas the aggressive reduction of power by several dB results in a much more pronounced artifact.

Power Related Artifacts: Low-Power Sawtooth Pattern

The output power level was found to fluctuate slightly at low-power settings (< 200 dB) with this causing subtle across-track banding artifacts that had a sawtooth-like nature (see Fig. 6). As with the case of the hysteresis problem, there was no solution available in post-processing, thus the recommendation to the survey crew was that power settings below 200 dB should not be used at all during main scheme survey in deeper waters (where, presumably there would less chance of saturation) and that power settings below 200 dB should only be used to avoid saturation in shallower waters.




Figure 6. Fluctuation in the power level for low power settings result in a slight variation in output level from ping to ping in a repeatable, sawtooth-like pattern as shown in the single pass mosaic of imagery acquired at a low power setting.

Both of these issues have since been addressed by R2Sonic with the strategy being to measure the actual transmit voltage on a ping-by-ping basis and to report the dB error associated with these two effects as a corrector to the nominal transmit voltage, both of these being reported in the header section of the ping datagram record. A Beta version of the firmware was available in time for the resumption of survey activities in Spring of 2014. The firmware was tested to ensure that the transmit voltage error corrector could be correctly applied in post-processing prior to upgrading the Substructure systems with this new firmware for their Spring operations. The corrections were verified (Figs. 7 and 8) and then implemented in QPS FMGT as of version 7.4.1 (released in July of 2014).




Figure 7. Example of power hysteresis effects and their corrections with and without the transmit voltage error applied.





Figure 8. Example of low-power sawtooth-like artifacts and their corrections with and without the transmit voltage error applied.

Improving Post-Processing

With the appropriate field procedures in place, the last portion of this effort was firstly to improve efficiency of backscatter processing and secondly to improve the quality of resulting imagery. Improving the efficiency was of prime importance for a number of reasons, the most important being that the scale of the operation necessarily leads to a very large amount of data to process. Bathymetric processing was already taking a significant amount of resources and the backscatter processing could impose very little additional requirements on the crew. QPS FMGT, which already provided a simple and straightforward workflow with data acquired using QPS QINSy (Fig. 9), was augmented to support dual-head R2Sonic import and very little else was required to streamline and/or simplify post-processing.

Stage 1: Import Data



Stage 2: Make Map




Figure 9. Schematic of the two-stage workflow involved with processing backscatter data in QPS FMGT. The tight integration with the QINSy qpd format allows FMGT to use the soundings computed in real-time for georeferencing of the backscatter data.

The second and most important part of the post-processing software improvements was to improve the quality of the output imagery since many in the community have found that commercial software has produced disappointing results with R2Sonic imagery. Seeing as the UNB SwathEd imagery has produced reasonable results since 2011 (Fig. 10), a thorough code review by Beaudoin and Doucet was undertaken in order to find and rectify the software errors that were causing the problems. The main source of error was the low resolution of the travel-time stored in the GSF record that was created during the DB/QPD merge process that occurred during data import. Other errors included a double removal of a travel-time correction that is applied by R2Sonic to correctly map a beam's travel-time back to the appropriate center sample and the double application of the sonar head dB reference corrector. Imagery mosaics showing the improvements can be seen in Fig. 11.




Figure 10. Example of imagery artifacts from initial R2Sonic processing efforts using commercial software as compared to the output from UNB SwathEd.




Figure 11. Comparison of mosaic output from UNB SwathEd and QPS FMGT after correction of software errors in FMGT.

Conclusion

As with any project, there still remains some work to do and the data processing from Substructure's Spring 2014 data set will provide a good test to confirm that the firmware upgrade and the associated correction procedures implemented in FMGT will consistently provide good quality imagery. The R2Sonic systems are still new to the market with respect to backscatter and there are still a few small remaining items to be addressed, for instance, deriving better default dB reference offsets for the various models.

There has been substantial and measureable progress in providing an R2Sonic workflow where none existed before, including acquisition best-practices, to the community. This happened in large part because of the close partnership between surveyors, researchers and hardware/software vendors involved. This model of close partnership has worked well and we hope to continue along these lines in the future to bring about further improvements.

References

- Multibeam Advisory Committee (MAC) (2013), "SVP Editor Software", available online at <http://mac.unols.org>
- Rice, G., Greenaway, S., Weber, T., and J. Beaudoin (2012), "Methods for Collecting and Using Backscatter Field Calibration Information for the Reson 7000 Series Multibeam", 2012 Canadian Hydrographic Conference, Niagara Falls, Canada, May 15 - May 17. Conference Proceeding

Lead author biography

Jonathan Beaudoin has a PhD (2010) in Geodesy and Geomatics Engineering from the University of New Brunswick. While he is employed by QPS in Fredericton, NB, Canada, he maintains an Affiliate Assistant Professor at the University of New Hampshire's Center for Coastal and Ocean Mapping, where he was employed as a Research Scientist from 2010-2013.