# Application of JPEG 2000 Wavelet Compression to Multibeam Echosounder Mid-water Acoustic Reflectivity Measurements

J. Beaudoin*

**Abstract**

The JPEG 2000 image compression standard is used to compress water column imagery data from Kongsberg Simrad multibeam echosounders (MBES) through integration of the JasPer software library into the Ocean Mapping Group (OMG) MBES processing software suite. An examination of data acquired over a sunken wreck demonstrates that lossy compression rates as high as 20:1 can be achieved with minimal signal distortion and little degradation of water column imagery quality. The data are also examined from an application specific viewpoint, that of wreck detection and measurement, with promising results even at compression levels exceeding 20:1. Though the methods used in this work were applied in post-processing, they are applicable in real-time. For exploration campaigns where sounding and water column imagery data are acquired continuously (e.g. CCGS Amundsen, NOAAS Okeanos Explorer), high compression rates may prove attractive as they would allow for acquisition of potentially useful water column information without the demanding storage requirements that currently limit the ability, or the desire, to record such data.

## 1 Introduction

In addition to measuring bathymetry and seafloor reflectivity, multibeam echosounders (MBES) can potentially record water column reflectivity measurements. These data have applications ranging from improved determination of the least-depth over wrecks [11, 12] to the detection of underwater gas plumes [9]. Figure 1 provides examples of water column reflectivity measurements being used to image the hull and mast of a sunken wreck. The reader is directed to [11] for more explanation of water column imaging geometry and interpretation. Despite the potential value of water column reflectivity measurements, the media storage requirements can be prohibitive and many MBES users do not routinely record water column reflectivities. A ten-fold increase in data storage requirements is not uncommon, this is especially problematic in shallow water applications as the higher repetition rates lead to data rates of several gigabytes per hour.

Data compression provides a potential solution to the problem of data storage, however, commonly used lossless data compression techniques yield only modest compression rates. For example, the Lempel-Ziv compression algorithm [18], which is used by many common desktop compression software packages, compresses the water column data examined in this work by approximately 25%, i.e. compressed data will be 75% of the uncompressed data size. Recent work has demonstrated that wavelet based lossy compression methods, in particular the JPEG 2000 standard, can be applied to acoustic signals with high compression ratios (20:1) and minimal signal distortion [6]. In the case of [6], seismic data (recorded in SEGY format) were converted to JPEG 2000 format for two reasons: (1) the wavelet compression techniques significantly reduced file sizes, and (2) the multi-resolution and random access capabilities of the JPEG 2000 standard increased responsiveness during digital access and improved accessibility to end-users. As will be shown, the approach taken by [6] is applicable, with slight modifications, for MBES water column imagery.

## 2 Background

### 2.1 Data Compression with Wavelets

Much like the Fourier transform uses a summation of sinusoidal functions to represent a signal, the wavelet transform allows for a representation of a signal as a superposition of wavelet basis functions. As wavelets basis functions can

*Ocean Mapping Group, University of New Brunswick, Fredericton, NB, Canada (now at Center for Coastal and Ocean Mapping, University of New Hampshire, Durham, NH, USA)

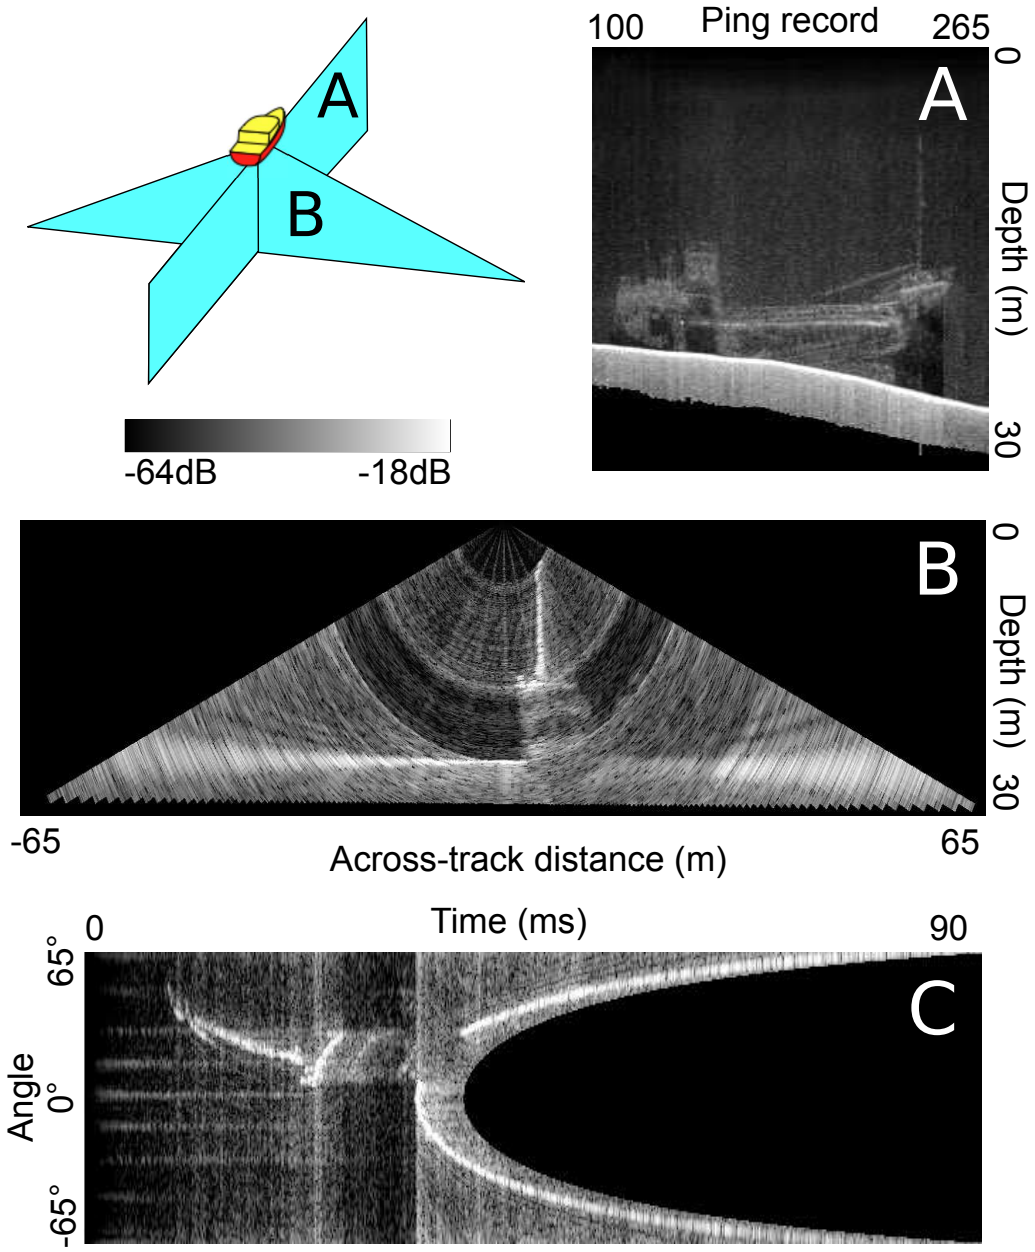


Figure 1: Sample water column imagery of a sunken wreck created by mapping acoustic backscattering strength (in decibels) to a greyscale image. The cartoon in the upper left depicts the imaging geometry associated with a MBES (in this case, a Kongsberg EM3002 [13]), with the triangular wedge B highlighting the portion of the water column and seafloor that is insonified during transmission of a single ping. In this example, the survey vessel passes over the longitudinal axis of the wreck from stern to bow. Examining the data from the central portion of the sector over several hundred pings collected along-track, it is possible to create a vertical section of the water column reflectivity as shown in Panel A. Panel A presents water column reflectivity in a manner similar to a single beam echosounder (SBES), however, the advantage of the wide mapping swath of a MBES is shown in Panel B in which it is clear that the survey vessel passed to the side of the wreck. Panel C shows a time-angle image from a single ping in which each row represents the return signal time-series for a receiver beam pointing in the direction associated with the angle on the y-axis. Angle is relative to the vertical with port side angles being negative. Note the curved appearance of the mast in the image, this is due to plotting polar coordinates (time and angle) in the Cartesian coordinate system of the image. In the Kongsberg recording format, data are only recorded for a short time beyond the bottom detection for each beam. Time-series records from inner beams are thus typically shorter than those for outer beams, this explains the apparent data gap for the majority of the right half of Panel C.

be localized in time, they allow for temporal discrimination of a signal's frequency content, this being their main advantage over non-windowed Fourier methods. This property allows for better representation of non-stationary signals and signals with discontinuities and/or sharp peaks [10].

Discrete orthogonal wavelet basis functions are generated through dilation and translation of a mother function

$$\psi^{a,b}(t) = |a|^{-\frac{1}{2}}\psi\left(\frac{t-b}{a}\right) \tag{1}$$

for a one-dimensional variable, $t$ with the variables $a$ and $b$ serving to dilate and translate the mother wavelet function [4]. The mother wavelet function must satisfy the condition

$$\int dx\ \psi(x) = 0 \tag{2}$$

which is to say that it must be oscillatory in nature [7]. It is also desirable for the mother function to decrease rapidly for $t \to \pm\infty$; this second requirement is necessary for the time/space localization properties of the wavelet transform [16].

The wavelet transformation of a signal is defined as the inner product of a function *f(t)* and a wavelet function $\psi^{a,b}(t) = |a|^{-\frac{1}{2}}\psi\left(\frac{t-b}{a}\right)$

$$(T^{wav}f)(a,b) = |a|^{-\frac{1}{2}}\int dt\ f(t)\ \psi\left(\frac{t-b}{a}\right) \tag{3}$$

in which the parameters *a,b* allow for tuning of frequency and space/time localization. A continuous wavelet transformation (CWT) results when both $a$ and $b$ are allowed to range over $\mathbb{R}$ (with the exception that $a \neq 0$). A discrete wavelet transformation (DWT), which is that most commonly used in image compression, results from restricting $a$ and $b$ to $\mathbb{Z}$ such that $a = a_0^m, b = nb_0a_0^m$, where $a_0 > 0, b_0 > 0, n \in \mathbb{Z}$. In this case, the transform is computed as

$$(T^{wav}_{m,n}f) = a_0^{-\frac{m}{2}}\int dt\ f(t)\ \psi(a_0^{-m}t - nb_0) \tag{4}$$

and the dependence of $b$ on $a$ interlocks the two parameters in such a manner that narrow wavelets are translated by small amounts and wider wavelets are translated by larger amounts. This ensures that the entire time (or spatial) range is covered during the transformation [7].

It is possible to chose $\psi$, $a$ and $b$ such that the $\psi_{a,b}$ provide an orthonormal basis for $L^2(\mathbb{Z})$, allowing for a multi-resolution analysis of a signal in which the signal is split into a pyramidal set of approximation signals where the resolution differs by a factor of two between each level. The difference between each level of resolution is encoded as coefficients of orthogonal "mother" wavelets that are dilated by factors of 2 (and translated by integer amounts) for each level of resolution [15]. In addition to the use of a mother wavelet to encode the incremental amount of information between resolution levels, a "father" wavelet, or scaling function, is used to approximate the signal with a downsampled version of the signal from the previous level. The downsampled approximation function is iteratively examined at each resolution level to yield (a) a further downsampled approximation signal, and (b) the incremental amount of detail information required to rebuild the previous resolution level's signal. The process may be repeated with the approximation signal being split into low and high frequency components until there is no more information to be extracted from the signal.

Retaining the mother wavelet coefficients throughout the iterative procedure yields as many wavelet coefficients as there were original data samples in the original signal (thus a multiresolution analysis is *O(n)* in terms of storage requirements). The DWT is implemented in practice through digital filters and the required number of operations at the first iteration of the pyramidal reduction is *O(n)*; the number of operations decreases by a factor of two at each stage due to downsampling, thus the total number of operations required is *O(2n)*.

A multiresolution analysis provides a wavelet transformation mechanism that is efficient in terms of processing speed and memory requirements. What makes this particular technique powerful for data compression is the fact

that the DWT of a signal can result in significant "energy compaction" [16] as the wavelet coefficients may allow for representation of a signal in a more redundant form in which many, or even the majority, of the wavelet coefficients may be of negligible amplitude. Figure 2 provides an example of the energy compaction properties of the wavelet transformation.

### 2.2 JPEG 2000 Image Compression Standard

The JPEG 2000 image compression standard was developed to improve upon deficiencies in the previous JPEG standard and to implement additional functionality. A significant difference between the previous standard and the new standard is the move from discrete cosine transformation (DCT) based coding to wavelet/subband coding techniques as in [4]. Coding can be done losslessly with biorthogonal 5/3 wavelets reversible transforms, as described in [5]. Lossless coding is implemented with a nonreversible and real-to-real 9/7 transform as proposed in [4]. After the wavelet transform, coefficients are quantized and encoded using bit-plane coding techniques and entropy coded with a binary arithmetic coder [1] as described by [17]. Further details regarding the standard and its implementation can be found in [1].

Software is available for encoding and decoding imagery using the JPEG 2000 standard, many open source applications use the JasPer software library [2]. This work relies on the JasPer library as well, in particular the functions which allow for encoding and decoding of data streams in computer memory (as opposed to conversion of imagery files written to non-volatile memory). Documentation for the JasPer software library can be found in [3].

## 3 Methods

### 3.1 Data

The data set examined in this work consists of water column imagery collected by a Kongsberg Simrad EM3002 MBES system by the CCGS Otter Bay in March-April of 2006. The EM3002 [13] operational frequency is approximately 300 kHz and can vary from this in the dual head mode of operation. The transmitter beam (single sector) covers an angular sector of approximately 130° with an along-track beam width of 1.5°. Receiver beams are 1.5° at broadside; as the receiver beams are electronically steered, the beam widths grow up to approximately 3° towards the edge of the covered sector. Though the system can provide 256 bottom detection solutions in its High Definition (HD) mode of operation, water column reflectivity measurements are only provided for 160 receiver channels. Data were acquired in a series of twelve passes over a 54 m long wreck in approximately 25 m water depth. Figure 3 shows a photograph of the vessel prior to being sunk in 1991, note that this is the same data set examined in [12]. The uppermost section of the mast was only fully imaged in one of the twelve passes over the wreck and this examination is limited to this single pass.

Raw water column data are stored as 8-bit signed characters and represent a time-series of echo levels in units proportional to decibels (8-bit signed integers represent a dynamic range of -64 to 63 dB, with 1 digital number = 0.5 dB). Note that positive backscattering values, though physically impossible, are accommodated in this representation to account for user-applied gains during acquisition. It should be noted that the time-series data are downsampled by the acquisition system to approximately the same range resolution as the pulse width [14].

### 3.2 Implementation

The imaging geometry associated with 2D seismic profiling (as in the work of [6]) is particularly suited to representation in a 2D image as there is a direct and intuitive logical mapping of 2D seismic data to a 2D image: each record (which consists of a time-series of backscattered echo intensities) is associated with a single downward looking beam, the data from which occupies a column in an image. As long as the digitizer sampling rate and the sampling interval (range) does not vary, then each record can be stored as a column in an image whose height would remain constant throughout the image. In the case of MBES water column data, each record consists of several time-series sub-records (one per physical receiver beam, for a total of 160 beams for an EM3002). One could follow the method of [6] and create a 2D image for each receiver channel (much like Panel A of Figure 1), however, the length of a given receiver beam's time-series can vary with depth over the course of a survey line (and so can the sampling rate). Furthermore, this approach does not lend itself to real-time compression.

In this work, images were created on a ping-by-ping basis by assembling the receiver sub-records into a time-angle image like Panel C of Figure 1 and then compressing the image using the JPEG 2000 standard. Pre-processing

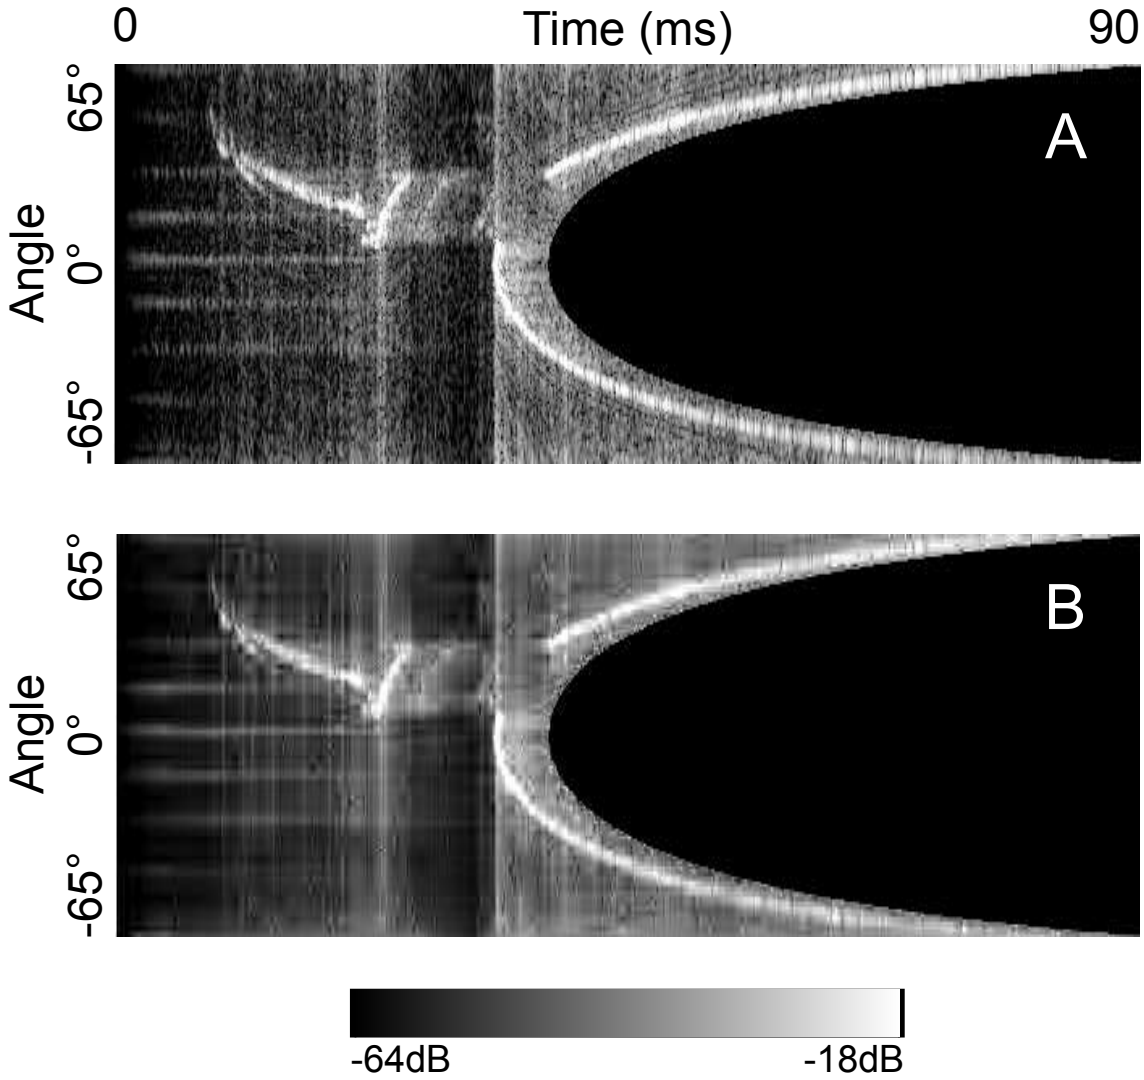


Figure 2: Sample time-angle water column imagery of a sunken wreck demonstrating the energy compaction properties associated with the wavelet transform using the Daubechies-4 wavelet [7]. In this example, a 2D DWT was applied to the raw time-angle water column image A above. The resulting wavelet coefficients were sorted in ascending order based on their absolute value. Coefficients below the 95th percentile were set to zero and the remaining large amplitude coefficients were used to reconstruct the image via an inverse DWT, resulting in the processed image B. Comparing the raw and processed imagery, it is clear that the energy compaction properties associated with the DWT allow for a reasonable reconstruction of water column imagery from only a small set (in this case 5%) of the largest amplitude wavelet coefficients. In other words, a significant portion of the wavelet coefficients are redundant; this can lead to significant gains in image coding efficiency during the quantization and entropy coding steps of wavelet image compression [8].

Figure 3: Photograph of the MV G.B. Church prior to being purposely sunk to create an artificial reef [12].

steps included joining the fragmented water column network packets into a single datagram [14], this is done during file format conversion between the Kongsberg and OMG file formats. Data from each beam are then inserted as a row in a 2D greyscale image in Portable Graymap (PGM) format. The width of the image is chosen to be large enough to accommodate the beam with the longest time-series; beams with shorter time-series are zero-padded. As the PGM format only supports unsigned integers, the signed 8-bit raw values are shifted positively by 128 to change the dynamic range to match that of an unsigned 8-bit integer (this includes zero-padded elements). Note that data remain logarithmically compressed at this stage.

Once the PGM formatted memory buffer is filled for a given ping, the JasPer library routines are used to convert it to JPEG 2000 format (JP2), the data are then written out to disk along with the Kongsberg Water Column datagram header [14]. Variable parameters at this point are limited to the choice of compression ratio (e.g. 10:1), though there are many other parameters to be investigated which may improve performance [3].

# 4 Results

Water column records were examined for a series of 170 swath records over the wreck, Figure 4 shows the raw image which is a subset of the image shown in Panel A of Figure 1. Lossless and lossy compressions were both attempted.

## 4.1 Lossless Compression

The primary result of interest for the lossless compression was the attainable compression factor. For the 170 pings over the wreck, achieved compression ratios ranged from 2.51:1 to 2.68:1 when comparing the file sizes associated with the PGM and JP2 formatted images. These are misleading results as the PGM image sizes are much larger than required to store the raw data due to the zero-padding involved during the pre-processing stage. The *effective* compression ratio, derived from comparing the JP2 formatted file to the raw data size prior to zero-padding, was much more modest with ratios ranging from 1.41:1 to 1.51:1 over the 170 pings. Though modest, the results are still an improvement over standard compression algorithms mentioned earlier: the lossless compression,

despite the penalty involved with zero-padding, packed file sizes down to 68% of the original file size on average as compared to 75% as achieved with the Lempel-Ziv compression algorithm [18]. It is expected that the effective compression ratio will grow with the angular sector of the system. For example, the increased angular sector of other Kongsberg sounders (e.g. EM710, EM302, EM122) would have an associated increase in the maximum time-series length relative to the shorter time-series at nadir. This would require additional zero-padding during the pre-processing stage where the time-angle image is constructed from the raw time-series data, leading to poorer effective compression ratios with lossless compression.

Improved performance can be achieved if one is willing to sacrifice data beyond the closest point of approach to the seafloor and limit the length of each receiver's time-series to the mininum of the set of beams in each ping. In this case, the compression ratios of 2.5:1 are readily achieved. For users interested in water column echoes only, this is a feasible option as the echogram is often contaminated by sidelobe interference beyond the closest point of approach to the seafloor. This option is not acceptable, however, for users interested in the echogram at or near the seabed across the whole swath.

## 4.2 Lossy Compression

The JPEG 2000 compression standard can be used for lossy compression. Evaluation of lossy compression of water column reflectivity data is difficult as there are many potential uses for such information. Hydrographers enjoy the improved confidence in least-depth determination over wrecks and other submerged hazards, thus they may prize geometric fidelity over radiometric fidelity. On the other hand, oceanographers imaging underwater hydro-thermal vents or water column processes such as internal waves may care more about image quality as the signal-to-noise ratio of their signal of interest can be much lower. Each application will likely have its own tolerance for the artifacts that can be introduced with lossy compression methods.

The JasPer implementation of the JPEG 2000 standard allows one to select a target compression rate, this being the reciprocal of a compression ratio (e.g. a compression rate of 0.1 corresponds to a compression ratio of 10:1). In this work, compression ratios ranging from 10:1 to 100:1 were applied with the goal being to ascertain at which point the lossy compression would introduce artifacts that might be considered unacceptable (this is inherently a subjective evaluation). A subset of these are shown in figures 5 through 8 (an uncompressed image is shown in Figure 4 for reference). It should be noted that the compression ratios reflect the compression ratio between the raw PGM size and compressed JP2 size. The *effective* compression ratio, which is computed by comparing the compressed JP2 image size to the total size of the raw data prior to zero padding, can differ by nearly a factor of two from the desired compression ratio. For example, a desired 10:1 compression ratio results in an effective compression ratio of 5:1. All compression ratios discussed below are desired compression ratios.

Examining the images of the wreck, it is clear that compression ratios of 10:1 and 20:1 yield imagery that is largely free of artifacts and suffers little loss of resolution. Most of the main deck rigging, masts and booms are still clearly identifiable and suffer little degradation of resolution. Higher compression ratios introduce more artifacts and degrade the quality of the imagery though the wreck is still identifiable as a wreck, even at 80:1 compression.

Quantitative metrics, such as correlation and RMS have been computed to provide an objective measure of lossy compression performance. Each of these metrics was computed for each swath with the sample-by-sample differences between the raw time-angle data and compressed time-angle data being examined as a set. These values were computed for each swath and then averaged over the set of 170 swaths in which the wreck was examined. The mean for each metric is plotted against the compression ratio in Figures 9 and 10. Both curves indicate that the majority of the loss in image fidelity occurs with compression ratios less than 20:1 with the gradient lessening for higher compression ratios. In other words, it is less costly to move from a 30:1 to 40:1 compression ratio than it is to move from 10:1 to 20:1. It is interesting to note that imagery remained nearly 100% correlated for a compression ratios of less than 5:1 (99.9941% correlated for 5:1). This low compression ratio could provide a reasonable alternative to the lossless method, which performed only marginally better than standard compression algorithms (recall that the desired ratio of 5:1 would result in an effective compression ratio of 2.5:1).

## 4.3 Application Specific Evaluation: Wreck Detection and Measurement

For hydrographic applications where safety of navigation is critical, real-time MBES bottom detection algorithms do not reliably track small and dispersed targets such as masts and rigging, especially when these targets are small with respect to the transmit and receive beamwidth of the MBES [12]. In these situations, water column reflectivity measurements can potentially provide the hydrographic surveyor with increased confidence that the least-depth

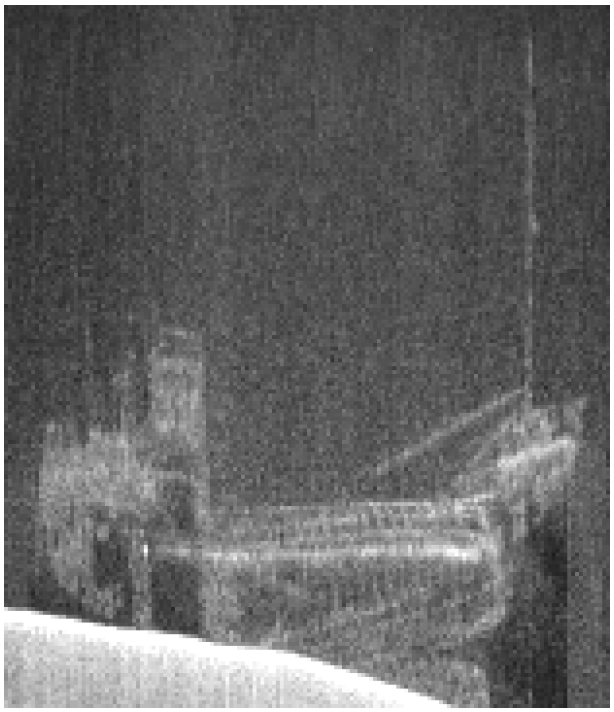

Figure 4: Raw water column image of sunken wreck (uncompressed).

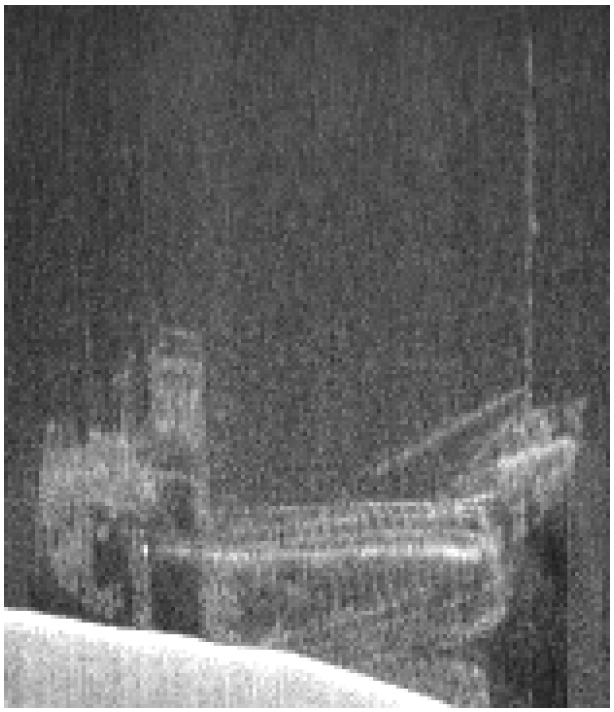

Figure 5: Compressed water column image of sunken wreck (10:1 compression ratio).

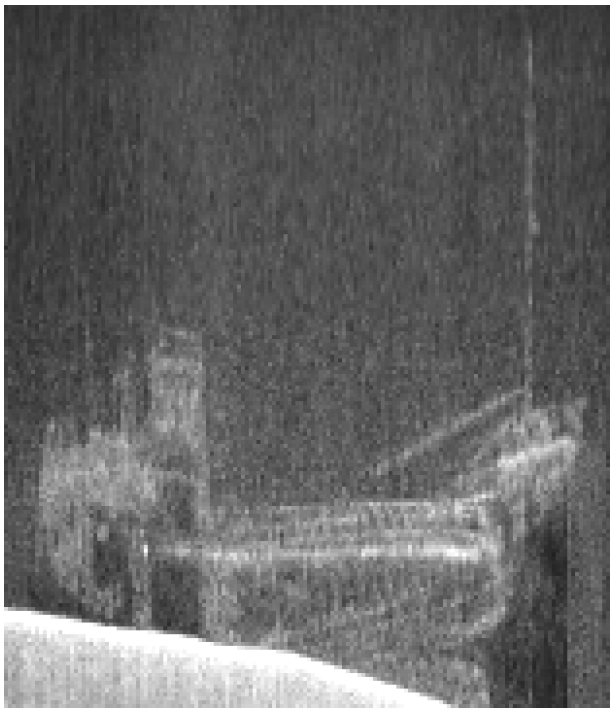

Figure 6: Compressed water column image of sunken wreck (20:1 compression ratio).

Figure 7: Compressed water column image of sunken wreck (40:1 compression ratio).

Figure 8: Compressed water column image of sunken wreck (80:1 compression ratio).

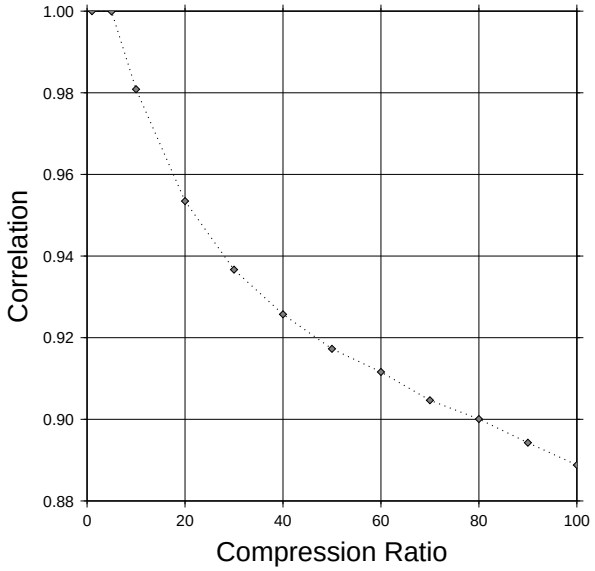


Figure 9: Mean image correlation plotted versus compression ratio.

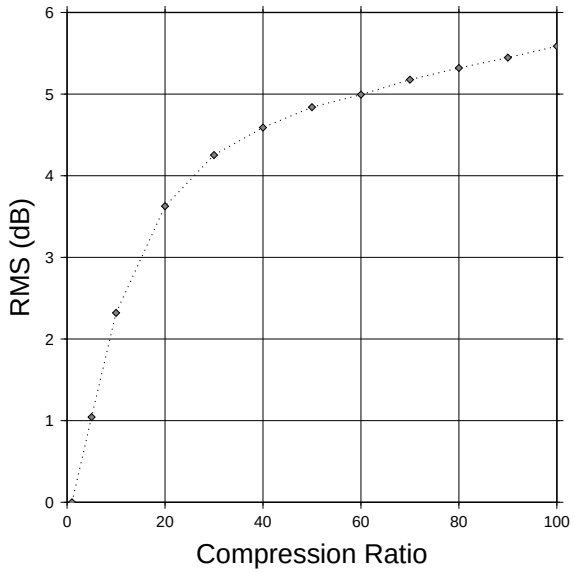


Figure 10: Mean RMS plotted versus compression ratio.

over sunken hazards has been detected. Lossy compression of water column data may, however, detract from the ability to (a) detect the hazard, and (b) to improve upon the least-depth estimates provided by the real-time bottom tracking algorithms.

As seen in the previous section, detecting hazards such as shipwrecks is still possible, even with high compression factors (see Figure 8). In this section, the impact of lossy compression on shipwreck and mast-detection is investigated using the wreck data examined thus far in this work, specifically the swath in which the fore mast had the highest signal-to-noise ratio (see panels B and C of Figure 1). The echo time-series from the 49 receiver beams that imaged the fore mast and upper deck of the vessel were examined, Figure 11 shows the echo time-series for the receiver beam that imaged the shoallest point of the mast along with the time-angle image of the wreck for ease of inter-comparison.

The ability to track mid-water targets with compressed data is assessed by first applying simple detection algorithms to the set of 49 receiver beam time-series from the raw data and comparing the results after applying the same detection algorithm to the data after compression has been applied. The discrepancies between the pre-compression and post-compression results serves as an application specific quantitative indicator of the impact of lossy compression, in this case, the ability to detect mid-water targets. For example, a simple thresholding technique can be used to detect targets in the receiver time-series. Figure 12 plots raw and compressed data from the same least-depth receive beam data shown in Figure 11. In this case, a simple threshold detection algorithm would appear to be relatively unaffected by lossy compression artifacts until very high compression rates are used (a threshold of -30dB was used). On the other hand, the change in slope and amplitude of the leading edge of the first return in the compressed data of Figure 12 makes a thresholding algorithm sensitive to the choice of threshold value. For this reason, three detection algorithms were tested to better understand how variations in a seemingly simple algorithm can make an application more or less susceptible to the artifacts associated with lossy compression. The three approaches are listed below.

- Thresholding: find the first sample whose amplitude exceeds a threshold of -30 dB
- Peak: find the peak return in the 50 sample window following the sample that first exceeds the threshold
- Weighted-Mean-Time: find the peak as in the last method, then compute the weighted mean time in a 50 sample window centered on the peak

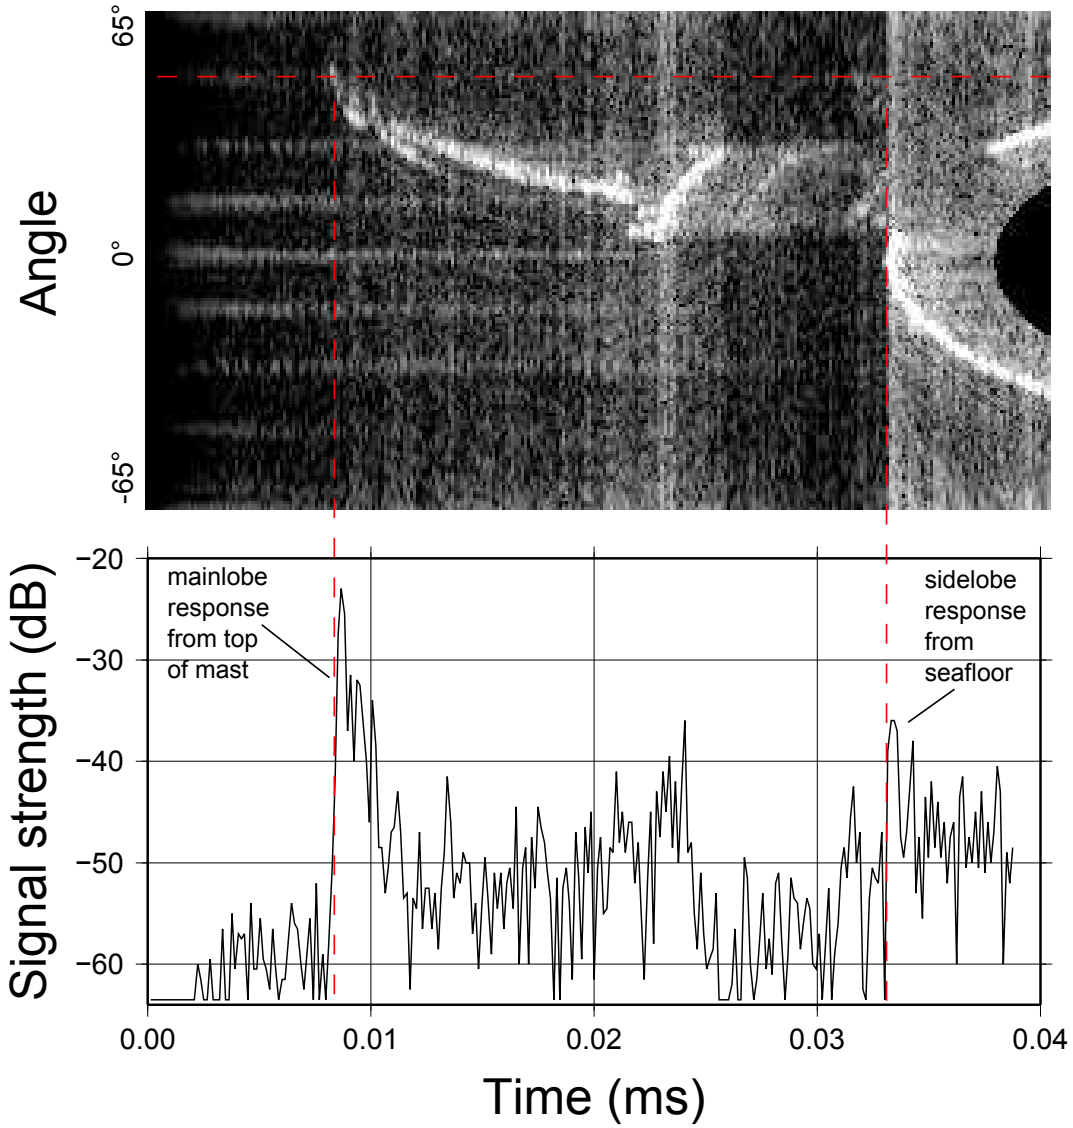


Figure 11: Time-angle image of the wreck (top) with the echo strength time-series from the top-of-mast receiver beam (bottom). The horizontal dashed line shows the location of the receiver beam's time-series in the time-angle image, vertical dashed lines point out the common reception times associated with various points of interest.

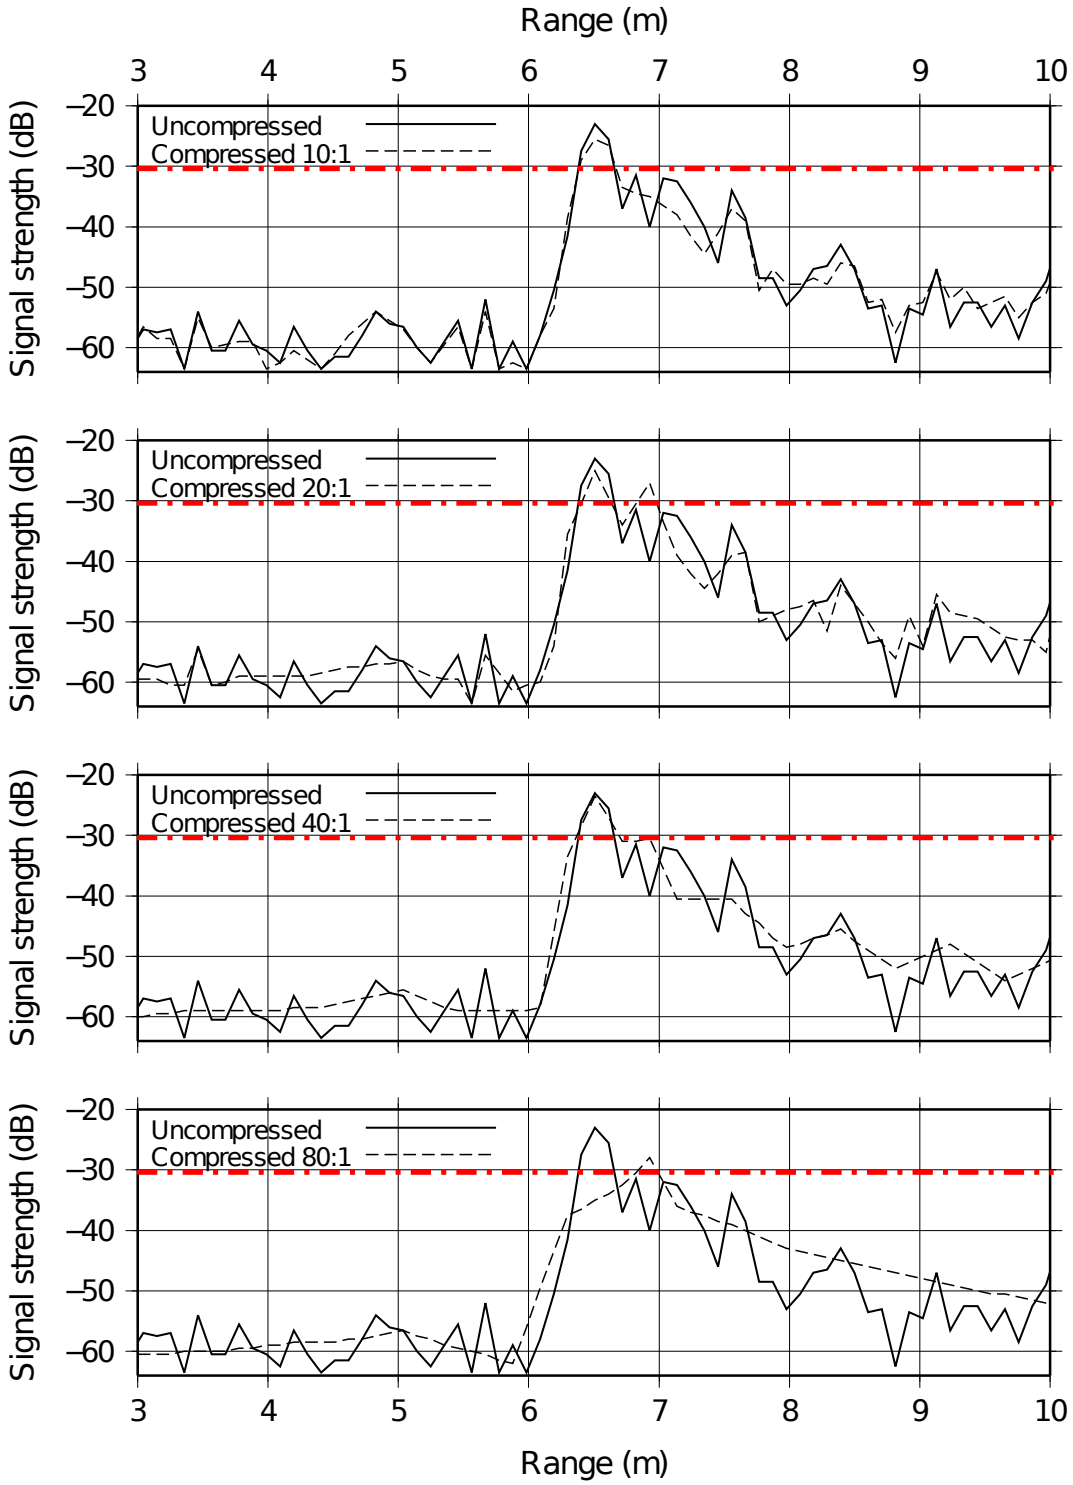


Figure 12: Time-series of received signal intensities for the top-of-mast receiver beam showing the effects of increasing levels of lossy compression. The raw (uncompressed) data are plotted as the solid line throughout the sequence of plots. Compressed data are plotted as dashed lines, with the compression factor varying by a factor of two between each plot (the subset of compression factors shown is the same as those used in the qualitative examination of the wreck, see Figures 5 to 8). Note the increasing loss of detail (and noise) with increasing compression rates. The red dash-dotted line indicates the signal threshold level of -30dB that is used in the various target detection algorithms.

The output from the raw detections consists of the range to the target, as detected in the raw time-series for each of the 49 receiver beams. Applying the same detection algorithm to compressed data yields a set of 49 range detections which differ slightly from the raw detections. Figure 13 plots the results of the Thresholding and Weighted-Mean-Time algorithms for the 49 receiver beams that imaged the upper deck and fore mast with varying levels of compression. Graphically evaluating the results from the figure alone, it is apparent that lossy compression artifacts distort the geometry of the deck and mast with increasing compression levels, however, the deck and mast are still successfully detected even at high compression levels.

Though the upper deck and fore mast are successfully detected for all compression ratios examined in Figure 13, the detections for the compressed time-series will invariably differ from the baseline detections, even if only by small amounts. The difference in detected range, referred to here as the range detection anomaly, is shown as a series of 2D histograms for each of the detection methods in figures 14 through 16. In these figures, the compression ratio ranges from 1 to 100 along the horizontal axis and each vertical column shows the distribution of range detection anomalies for the 49 receiver beams that imaged the mast and deck. Positive values indicate an overshoot in the detected range, i.e. over-predicting the range to the mast and/or upper deck. The mean and standard deviation, both as a function of compression factor, are plotted in figures 17 and 18 for all three 2D histograms.

The Thresholding and Peak algorithm detections can vary from the baseline detections significantly, even at low compression levels and applications that rely on these methods should proceed with caution or at least examine the sensitivity of their application to lossy compression. Spurious mistrackings of 1 m or more can occur with small targets and/or high sidelobe responses at or near the threshold level. An example of this is a mid water target tracked in Figure 13(b), located approximately 3.5 m to the right and 4 m above the lower left grid corner. The target is partially detected in the 10:1 data but is filtered in the images that have undergone higher compression factors and is thus not detected. This leads to a systematic range detection anomaly for the few beams that imaged the mid-water target in the raw imagery, these appear as positive outliers of 1.3 m in the 2D histogram of Figure 14. In this case, the raw detections are likely due to a sidelobe response from the mast and the compressed data are thus more correct in the sense that they are tracking the mainlobe response of the deck and not a mid-water sidelobe response from the mast.

Compared to the first two algorithms, the Weighted-Mean-Time algorithm appears to be much more resilient to the effects of lossy compression artifacts. Though the edges, slopes, highs and lows of the signals may become distorted by lossy compression, the center of mass of the signals remains relatively unaffected to the point that the Weighted-Mean-Time algorithm consistently outperforms the other two algorithms. Furthermore, the performance of the algorithm worsens only slightly with increasing compression. For the purposes of range measurement, this type of algorithm would appear to be the most applicable when using lossy compression methods. The sophistication of the other two algorithms could have been improved to yield better or comparable results, but the point of this exercise was not to develop perfect detection algorithms or to judge the best of the three. Instead, the point was to examine the circumstances under which common approaches to simple signal processing problems could be hindered by the introduction of artifacts due to lossy compression. In this case, it is evident that some approaches to target detection are more susceptible to uncertainty than others after undergoing signal compression, however, their performances are remarkably good considering the gains that are made in addressing the issue of storage requirements for water column data: the uncertainty in target detection and range measurement is a small price to pay for lossy compression methods that would allow for near continuous logging of this type of data.

# 5 Conclusion

The problem of compressing water column imagery from a MBES was examined using wavelet/subband coding techniques as implemented by the JPEG 2000 image compression standard, more specifically the JasPer implementation of the standard. Lossless methods provided a modest improvement over traditional compression algorithms, yielding effective compression ratios of approximately 1.5:1. Lossy methods proved more promising with effective compression ratios of 5:1 and 10:1 yielding very little loss of resolution or introduction of artifacts. It has been demonstrated that target recognition and detection is still possible, even with high compression ratios. Caution should be exercised however when applying lossy compression methods to water column imagery data as application specific objectives may vary in their tolerance of artifacts that are introduced with the use of lossy compression methods. Testing procedures should be implemented on an application-by-application basis to help objectively determine a balance between compression rate and signal distortion.

For exploration campaigns where sounding and water column imagery data are acquired opportunistically (e.g.

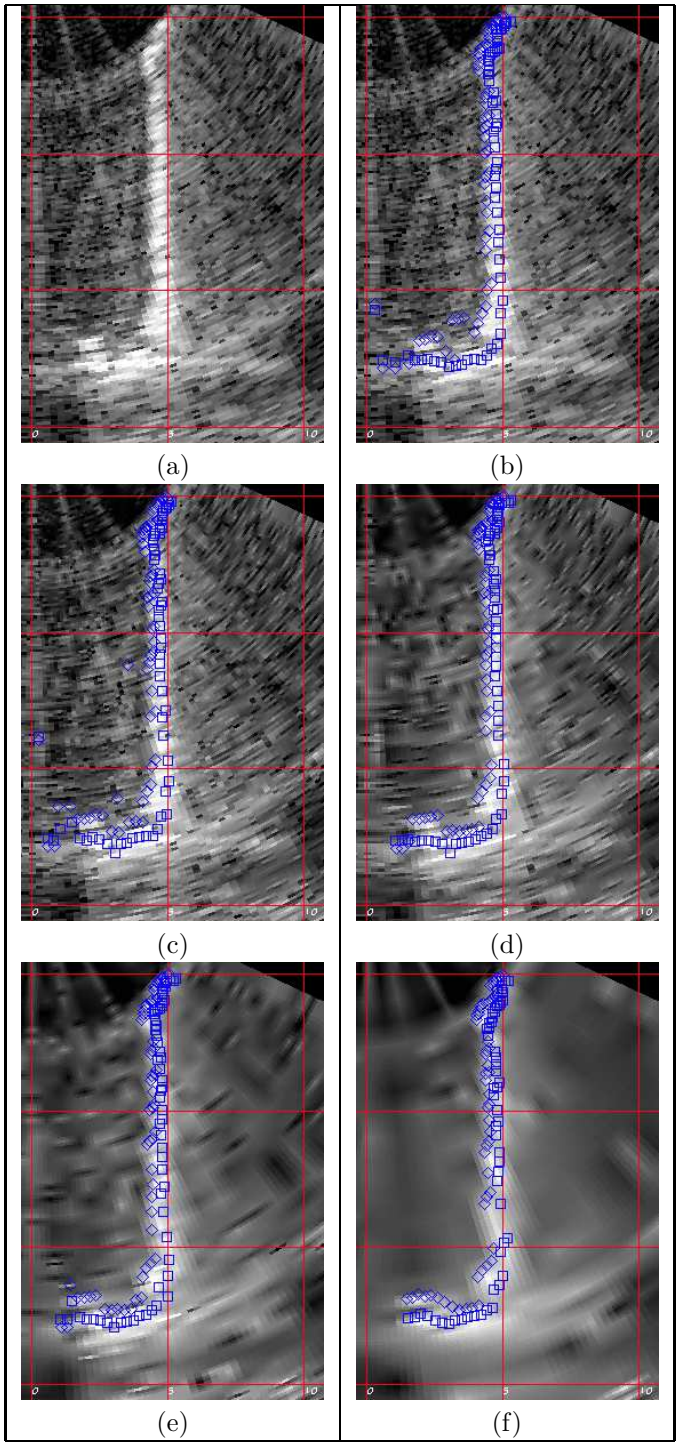


Figure 13: Subsets of Figure 1(b) showing raw (a and b) and compressed (c through f) imagery of the mast and deck with results from the Thresholding and Weighted-Mean-Time target detection algorithms plotted as blue diamonds and blue squares, respectively (results from the Peak detection are omitted for clarity). Image (a) provides an uncluttered reference image of the uncompressed data with image (b) being the same but with target detection results plotted as well. Images (c), (d), (e) and (f) show the same image after application of compression, with ratios ranging from 10:1, 20:1, 40:1 and 80:1, respectively. Target detection points in these four images result from applying the detection algorithms to the compressed time-series data stored in the rows of the JP2 image files. Note that the deck and mast is detected even at the highest compression level (80:1, image f) despite the poor quality of the imagery. Red grid lines are plotted at 5m intervals for scale.

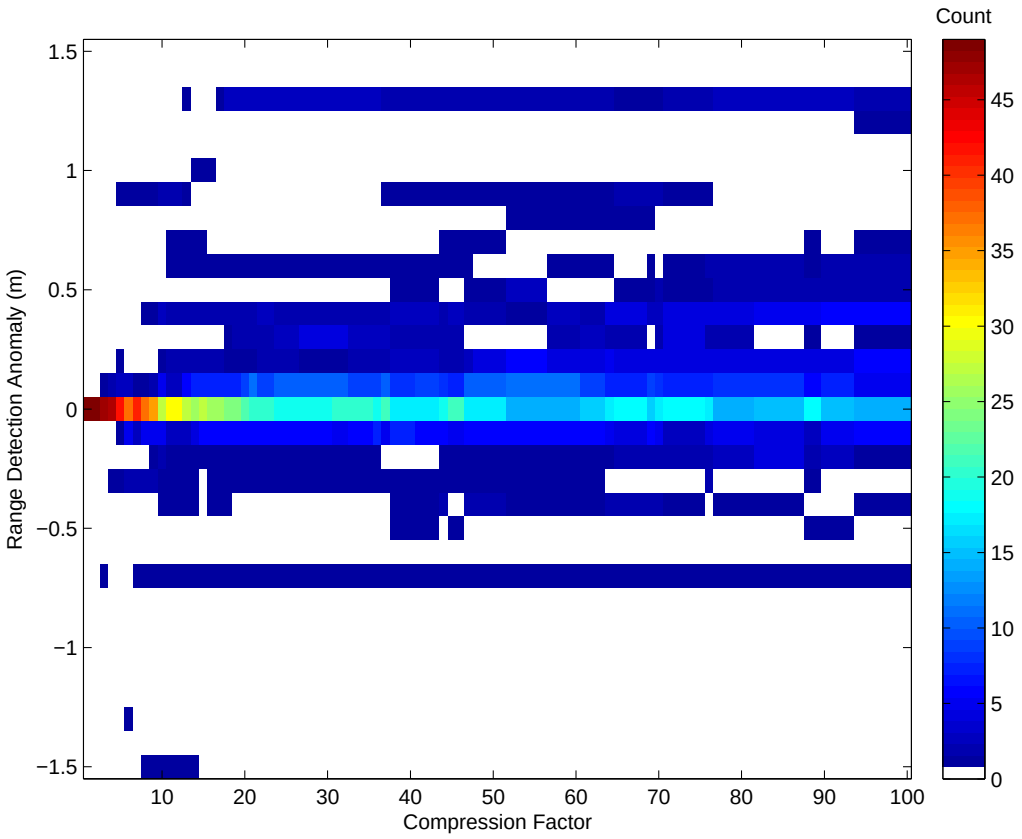


Figure 14: 2D histogram of range detection anomalies for the Thresholding target detection algorithm.

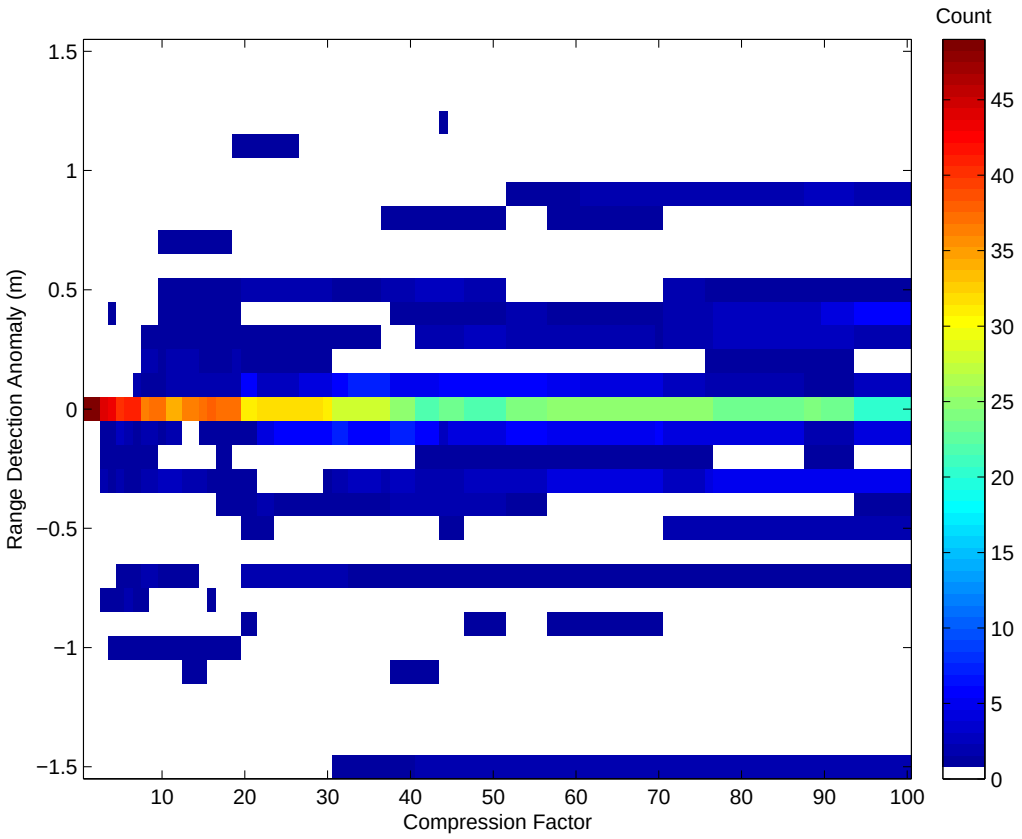


Figure 15: 2D histogram of range detection anomalies for the Peak target detection algorithm.

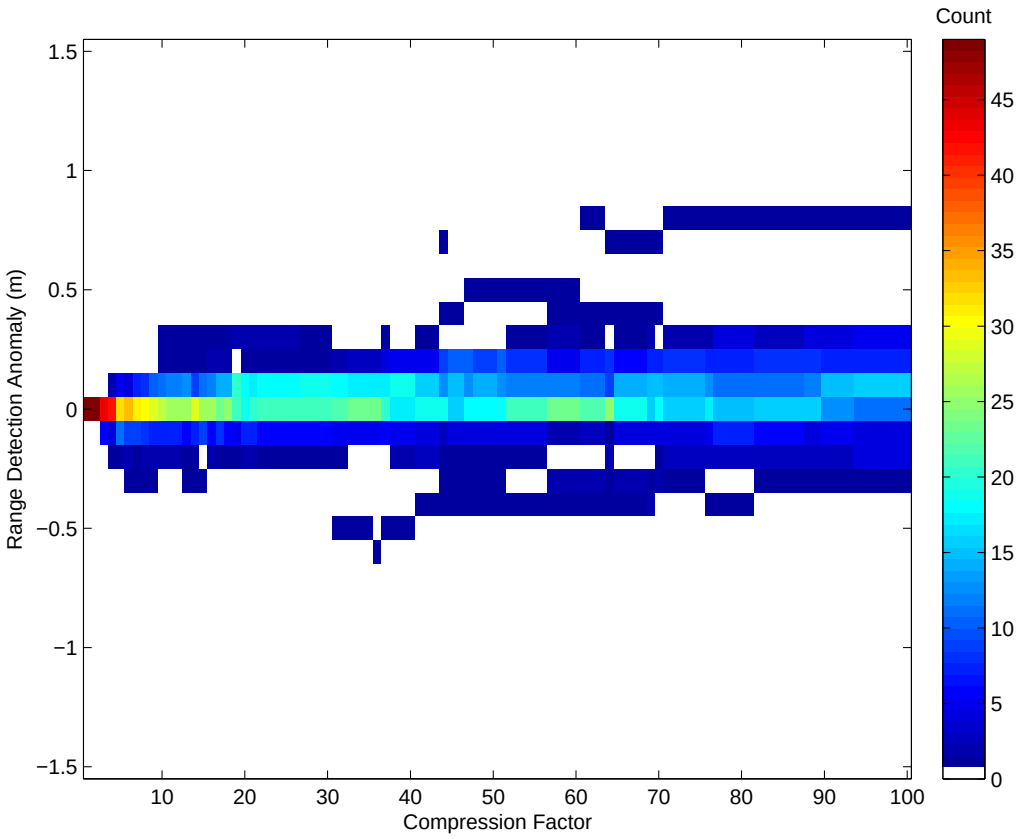


Figure 16: 2D histogram of range detection anomalies for the Weighted-Mean-Time target detection algorithm.

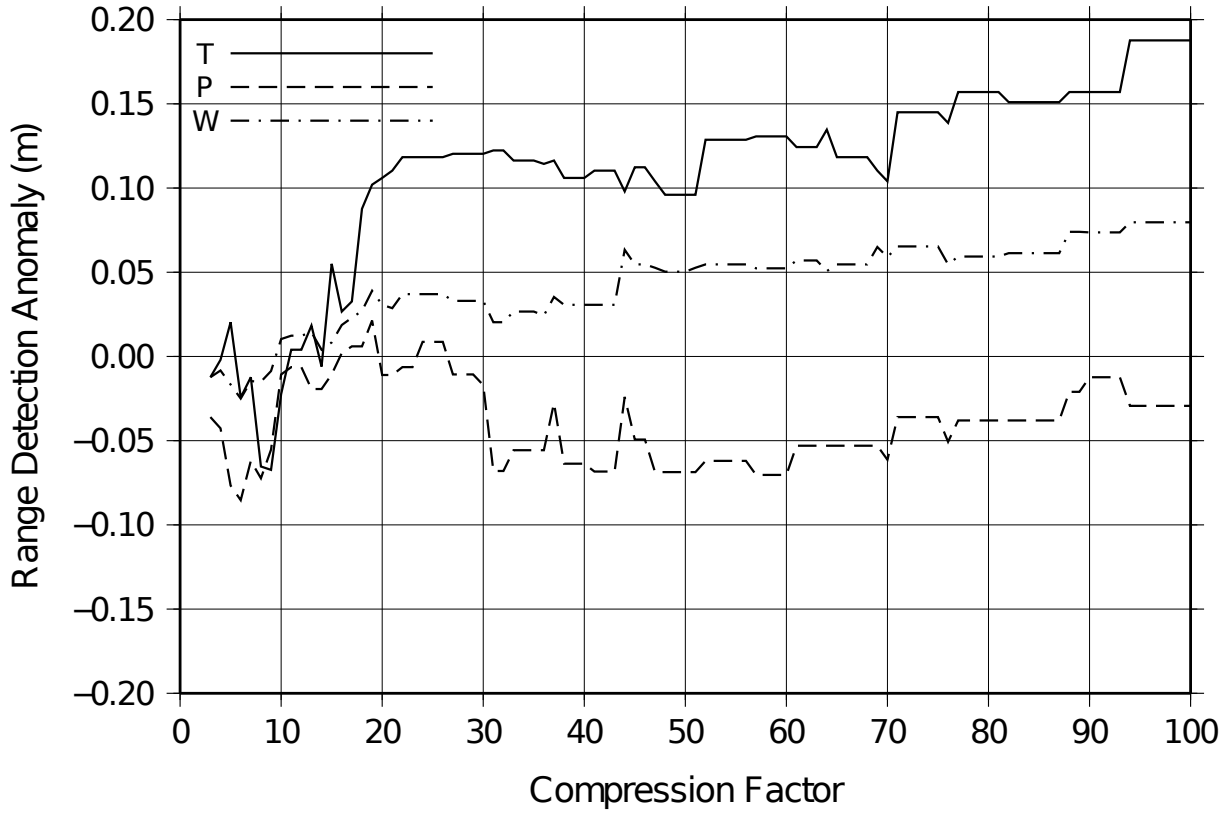


Figure 17: Mean range detection anomalies for the Threshold (T), Peak (P) and Weighted-Mean-Time (W) detection algorithms as a function of compression factor. The water column datagram data are downsampled to approximately 10 cm resolution; the mean detection anomaly is less than 1 sample on average for all three detection methods.

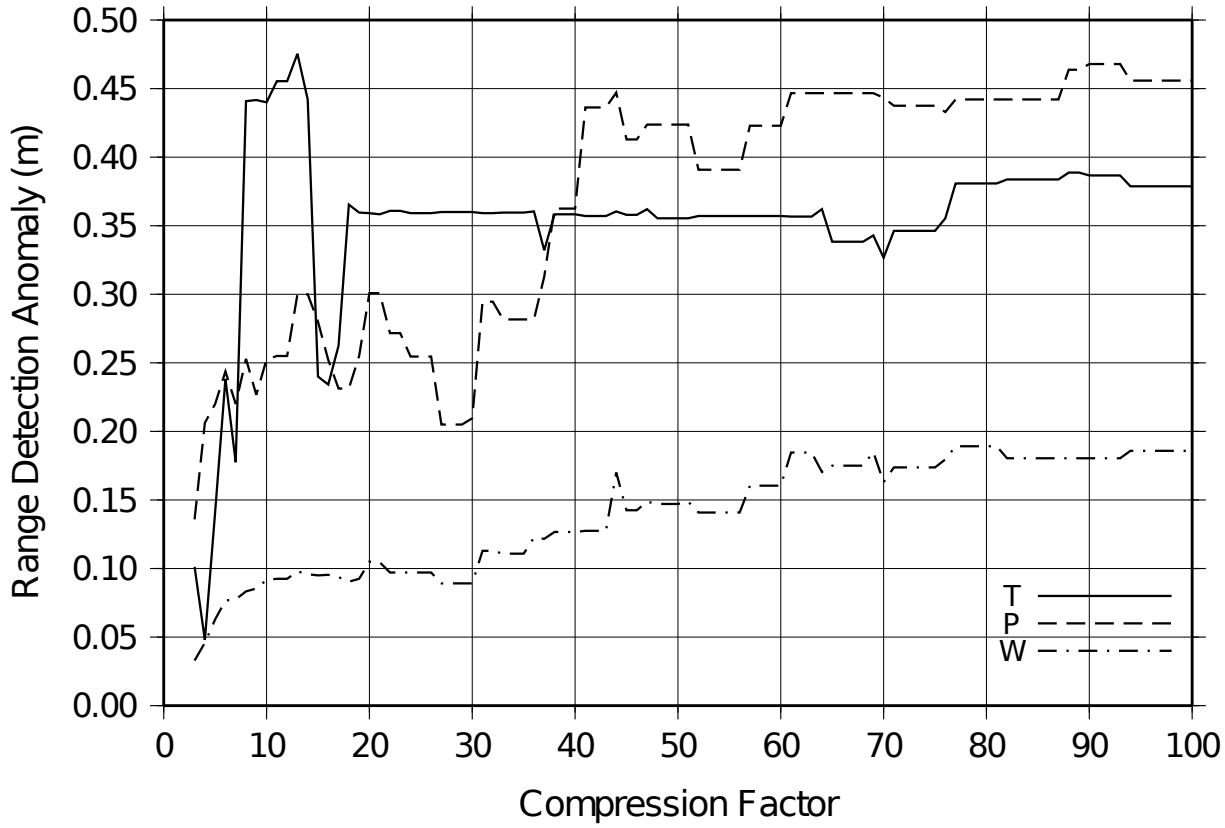


Figure 18: Standard deviation (1-sigma) of range detection anomalies for the Threshold (T), Peak (P) and Weighted-Mean-Time (W) detection algorithms as a function of compression factor. Recall that the water column data are downsampled to 10 cm resolution; thus the detection methods introduce 1 to 5 samples of uncertainty (1-sigma).

CCGS Amundsen, NOAAS Okeanos Explorer), high compression rates may be attractive as they would permit acquisition of potentially useful water column information without the associated storage requirements. Much like a sounding data from an MBES, the time series of reflectivities from a single receiver beam is seldom examined in isolation, nor is the image from a single swath examined without the context of neighbouring swaths. Data from neighbouring receiver beams and swaths could provide context that may help the human interpreter see through artifacts associated with high compression rates.

# 6 Future Work

Continual improvement, refinement and testing of the data model is foreseen in the near future, specifically investigation of ideal data representation (logarithmic or linear) and pre-preprocessing of the data (e.g. inter-sector normalization for multi-sector systems). Alternate data organization methods, e.g. tiling, will be investigated to aid in reducing the performance penalties associated with zero-padding the time-angle imagery. Once stable, additional efforts will focus on the integration of these compression procedures into acquisition and/or processing software (this has already been done for the OMG SwathEd software suite).

# 7 Acknowledgments

This work has been made possible by the sponsors of the Ocean Mapping Group and by the ArcticNet Network of Centres of Excellence. Many thanks go the officers and crew of the CCGS Otter Bay for their aid in acquisition of the data set assessed in this work. Finally, the author wishes to thank Robert Courtney of the Geological Survey of Canada (Atlantic). His early work with the application of the JPEG 2000 compression standard to sidescan and seismic data has inspired this endeavour.

# References

[1] ISO/IEC 15444-1: Information technology - JPEG 2000 image coding system - Part 1: Core coding system. 2000.

[2] M. Adams. JasPer: A software-based JPEG-2000 codec implementation. In *Proc. of IEEE International Conference on Image Processing*, Oct 2000.

[3] M. Adams. *JasPer software reference manual - ISO/IEC JTC 1/SC 29/WG 1 N 2415*. Dec 2002.

[4] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies. Image Coding Using Wavelet Transform. *IEEE Transactions on Image Processing*, 1(2):205–220, 1992.

[5] A. Calderbank, I. Daubechies, W. Sweldens, and B. Yeo. Wavelet transforms that map integers to integers. *Applied and Computational Harmonic Analysis*, 5(3):332–369, 1998.

[6] R. Courtney. Storage and dissemination of segy data in JPEG2000 format. In *Proc. Shallow Survey 2008*, Oct 2008.

[7] I. Daubechies. *Ten Lectures on Wavelets*. SIAM, Philadelphia, 1992.

[8] G. M. Davis and A. Nosratinia. Wavelet-based Image Coding: An Overview. *Applied and Computational Control, Signals and Circuits*, 1:25–48, 2009.

[9] J. V. Gardner, M. Malik, and S. Walker. Plume 1400 Meters High Discovered at the Seafloor off the Northern California Margin. *Eos Trans. AGU*, 90(32), 2009.

[10] A. Graps. An Introduction to Wavelets. *IEEE Computational Sciences and Engineering*, 2(2):50–61, 1995.

[11] J. E. Hughes Clarke. Applications of multibeam water column imaging for hydrographic survey. *The Hydrographic Journal*, April 2006.

[12] J. E. Hughes Clarke, M. Lamplugh, and K. Czotter. Multibeam water column imaging: Improved wreck least-depth determination. In *Proc. Canadian Hydrographic Conference 2006*, May 2006.

[13] Kongsberg Maritime AS. *EM3002 Multibeam Sonar Operators Manual*. 2004.

[14] Kongsberg Maritime AS. *EM Series Datagram formats – Format Descriptions*. Jan 2006.

[15] S. Mallat. Multiresolution Approximations and Wavelet Orthonormal Bases of L2(R). *Transactions of the American Mathematical Society*, 315(1):69–88, 1989.

[16] H. Stark. *Wavelets and Signal Processing – An Application-Based Introduction*. Springer-Verlag, Berlin, 2005.

[17] I. Witten, R. Neal, and J. Cleary. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–540, 1987.

[18] J. Ziv and A. Lempel. A Universal Algorithm for Sequential Data Compression. In *IEEE Transactions on Information Theory*, volume 23, pages 337–343, May 1977.