Processing C-PROOF glider Sea-Bird GPCTD data
June 14, 2023
H. V. Dosser

Version history:
Version 0.0 - October 7, 2022 - original document
Version 0.1 - February 8, 2023 - expanded the Introduction, added sections on removal of anomalous conductivity values and the sensor alignment correction, and updated thermal lag correction section.
Version 0.2 - June 4, 2023 - updated the figures, changed the example mission, added alignment section.

Preamble: This document describes conductivity, temperature, and pressure data processing procedures for delayed-mode data collected using Sea-Bird Scientific Glider Payload Conductivity Temperature Depth (GPCTD) sensors mounted on C-PROOF autonomous ocean gliders. This document provides as an example of the corrections applied to the **Calvert Line** mission **dfo-mike579-20210704**.

Table of contents:

1 Introduction and GPCTD sensor overview

Gliders in the C-PROOF fleet are equipped with one of two types of sensors to measure conductivity, temperature, and pressure in the ocean: Sea-Bird GPCTDs and RBRlegato CTDs. The Sea-Bird CTD is pumped, whereas the RBR CTD is unpumped. This document describes processing steps applied to delayed-mode data from the **pumped Sea-Bird GPCTD sensors**. This sensor has a maximum sampling rate of 1 Hz and was designed specifically for Slocum gliders.

Figure 1.1 shows the GPCTD dimensions, and the location of the water intake and exhaust. According to the Sea-Bird User Manual (2021) (emphasis added):

“The T-C sensor assembly visible on the exterior of the vehicle consists of a streamlined T-C intake sail (with integral T-C duct and anti-foul device), a horizontal, internal field conductivity cell, and a downstream exhaust sail with pump connections. The intake sail allows measurements to be made outside of the vehicle’s boundary flow where old water is thermally contaminated by the vehicle, producing TS errors. **The pump pulls water into the duct at top of the intake sail and immediately past a temperature sensor. Water then flows** through an anti-foulant cylinder, **through the conductivity cell**, and out the top of the exhaust sail to prevent exhaust re-circulation and Bernoulli pressure differences from changing the flow rate. The outside of the conductivity cell is free-flushed to minimize salinity errors. If the cell were located inside the flooded fairing, **a thermal mass error** resulting from temperature difference between the poorly-flushed volume inside the hull and the ambient ocean temperature measured by the CTD would produce salinity errors.”

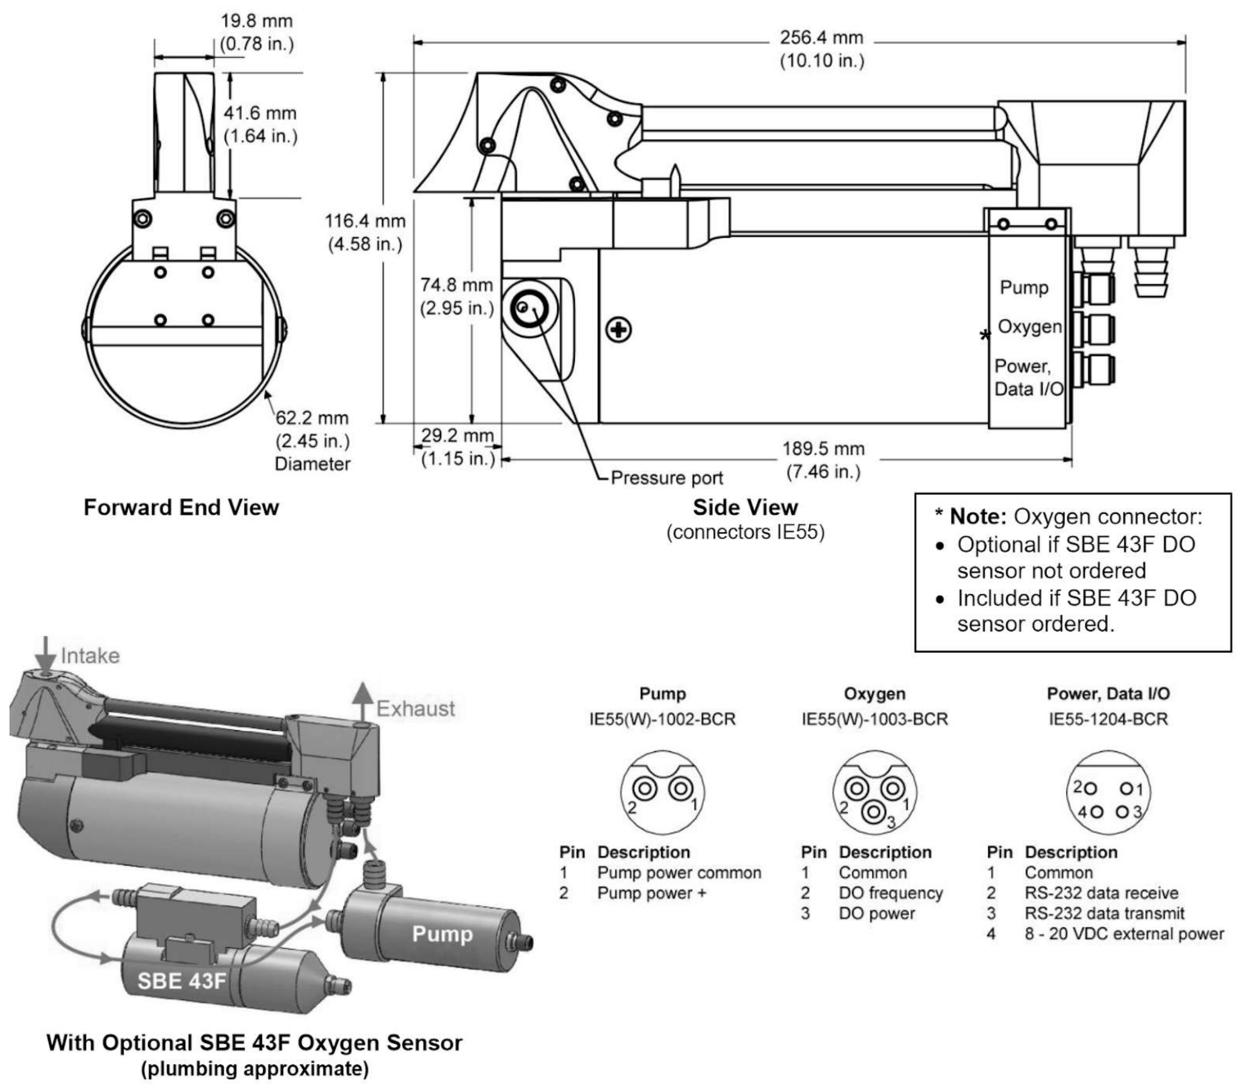

Figure 1.1: Diagram showing dimension and connectors for the Sea-Bird GPCTD (Sea-Bird User Manual, 2021). Note that C-PROOF gliders do not use the optional SBE 43F oxygen sensor.

The GPCTD has a maximum sampling rate of 1 Hz and a pumped flow-rate through the sensors of 10 ml/s. Water passing into the instrument will reach the temperature sensor almost immediately, then travel through the conductivity cell, the outside of which is in contact with the external environment. The pressure port is located below the intake on the side of the instrument. Each sensor has a different response rate, a different thermal inertia, and is physically separated from the other sensors, all of which results in the need for corrections to the temperature and conductivity data.

2 Corrections applied to delayed mode data

The two main issues commonly observed in the derived salinity data from these sensors are spikes in the salinity profiles, and an offset between subsequent downward and upward salinity profiles. The primary causes of these issues are 1) mis-alignment of temperature and conductivity measurements due to differing sensor response times and the physical separation between the thermistor and the conductivity sensor causing a transit time delay for water being pumped through the CTD and 2) the thermal lag effect caused by the thermal inertia of the conductivity cell affecting the temperature and thus conductivity of the water as it passes through the cell. The temperature sensor also has some small thermal inertia, however this is minimized by the design of the GPCTD and a correction for the thermal lag effect for the temperature sensor is not required. This document outlines two corrections designed to address these issues in the derived salinity data: a sensor alignment correction and a thermal lag correction (also referred to as a thermal mass correction) for the conductivity sensor.

In addition to applying these two corrections with constants confirmed or adjusted to be appropriate for our specific sensors, we also remove unphysical conductivity values and questionable salinity profiles (the latter of which is referred to as a 'wild edit' in the Sea-Bird processing manual). The unphysical conductivity values are removed before the sensor alignment correction is applied and the questionable salinity profiles are removed before the thermal lag correction is applied.

2.1 Identification and removal of unphysical conductivity values

2.1.1 Description of the correction to be applied

We first identify any conductivity values that are obviously unphysical, which is typically caused by air bubbles in the conductivity cell. We identify such values using a simple criterion applied to the raw conductivity data. The data is first binned by profile index, in increments of typically 100 profiles. The use of profile bins rather than depth or temperature bins is designed to allow for the removal of unphysical values in different regions samples during the same mission (e.g., shelf and open ocean) where the distributions of conductivity can differ significantly. The criterion then temporarily flags any data points that are more than a specified number of standard deviations (typically 5) away from the mean value for each bin, then recomputes the mean and standard deviation, excluding the temporarily flagged values. Conductivity values that still differ from the mean by more than a specified number of standard deviations (typically 5) are flagged as 'bad' and set to NaN. If the difference between the 'bad' values and the mean is less than the accuracy of the sensor, which is 0.0003 S/m for the GPCTD, then those points are not excluded. The number of profiles per bin may be adjusted, as may the number of standard deviations.

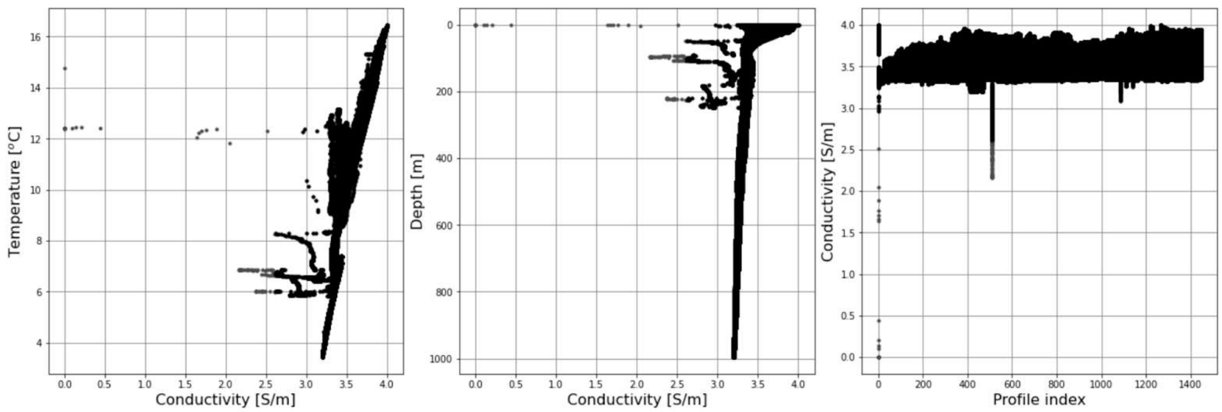

Figure 2.1: Temperature vs. conductivity (left panel), depth vs. conductivity (middle panel), and conductivity vs. profile index (right panel), with the red dots indicating the unphysical values that have been flagged as bad and set to NaN.

2.1.2 Application of the correction to mission dfo-mike579-20210704

This mission on the Calvert Line began and ended in Hakai Pass, so that the glider was deployed and recovered close to shore in a coastal region where the distribution of conductivity values differed significantly from those sampled while crossing the continental shelf in Queen Charlotte Sound. The data were binned in increments of 100 profiles, and 5 standard deviations were used for the criterion to flag and exclude unphysical conductivity values.

For this mission, we flag and exclude the extremely low conductivity values occurring at the surface (Figure 2.1) consistent with air bubbles in the cell, as well as some spikes in conductivity deeper in the water column. Some unphysical values are missed by this correction, including portions of these sudden drops in conductivity between the surface and 250 m, and may be caught during the removal of unphysical salinity profiles below.

2.2 Sensor alignment correction

2.2.1 Description of the correction of be applied

The sensor alignment correction is used to align conductivity (C) with temperature (T) in time. This correction reduces the occurrence of salinity spikes near sharp gradients in T and C and ensures salinity calculations are made using the same parcel of water. In practice, this alignment correction is made to both T and C relative to pressure.

Janzen and Creed (2011) report manufacturer-recommended values to use for both the sensor alignment and thermal lag corrections for the Sea-Bird GPCTD. They state that "Alignment corrections were based on the transit time between the T-C sensors, the temperature response time, and the estimated response time of the conductivity sensor in a 10 ml/s flow." The following constants are used for their Sea-Bird GPCTD alignment corrections:

- Temperature advanced by +0.5 s relative to pressure
- Conductivity advanced by +0.4 s relative to pressure

Our procedure for the sensor alignment correction is as follows:
Step 1: Directly estimate the sensor alignment correction constant for conductivity relative to temperature, τ_C following the methodology of Ferrari and Rudnick (2000).
Step 2: Confirm that the directly estimated value of τ_C is within ±0.5 s of the Janzen and Creed (2011) value of τ_C = -0.1 s (and if a given sensor has a directly estimated value of τ_C that is consistently higher or lower than -0.1 s, investigate further).
Step 3: Apply the alignment correction using the alignment constants for conductivity and temperature provided in Janzen and Creed (2011), and specified above.

We directly estimate the sensor alignment correction constant using the following procedure:

a. Assume conductivity is approximately linearly related to temperature
b. Calculate the cross-spectrum between the standardized temperature and conductivity timeseries
c. Calculate the coherence, to determine where the cross-spectrum is valid
d. Fit the function $2\pi f \tau_C$ to the phase of the cross-spectrum, where f is frequency and τ_C is the correction constant we are looking for

These steps are based on the methodology of Ferrari and Rudnick (2000) and applied to each individual profile, with the highly variable surface layer excluded based on a cutoff density chosen by examining a T-S diagram. This cutoff density is chosen to be above the main thermocline.

The spectrogram for the cross-spectra from individual profiles is determined, where the cross-spectrum is estimated using Welch's method, with a Hann window of length 128 and 50% overlap. Values are only considered valid if the squared coherence is at least 75% and the 95% confidence level is exceeded. The choice of 75% is arbitrary, and results are not sensitive to this choice within the range 50% to 90%.

We fit the function $2\pi f \tau_C$ to the phase of the cross-spectrum from all profiles, with a frequency cutoff chosen based on a drop in mean coherence and a sign change in the slope of the mean cross-spectrum phase. Provided that the resulting value for the time constant τ_C is within the required ±0.5 s of the value specified for Sea-Bird GPCTDs, τ_C = -0.1 s, we confirm that the Janzen and Creed (2011) constants for the alignment correction are appropriate for use with the GPCTD in question. Note that we preferentially use the Janzen and Creed (2011) constants over our directly estimated value for τ_C since we have found that the direct estimate is sensitive to the choices of frequency cutoff and density cutoff used, with reasonable choices producing constants that can vary by up to ±0.5 s. Finally, we apply the correction using a linear time interpolation method, based on the Sea-Bird CTD data processing manual.

2.2.2 Application of the correction to mission dfo-mike579-20210704

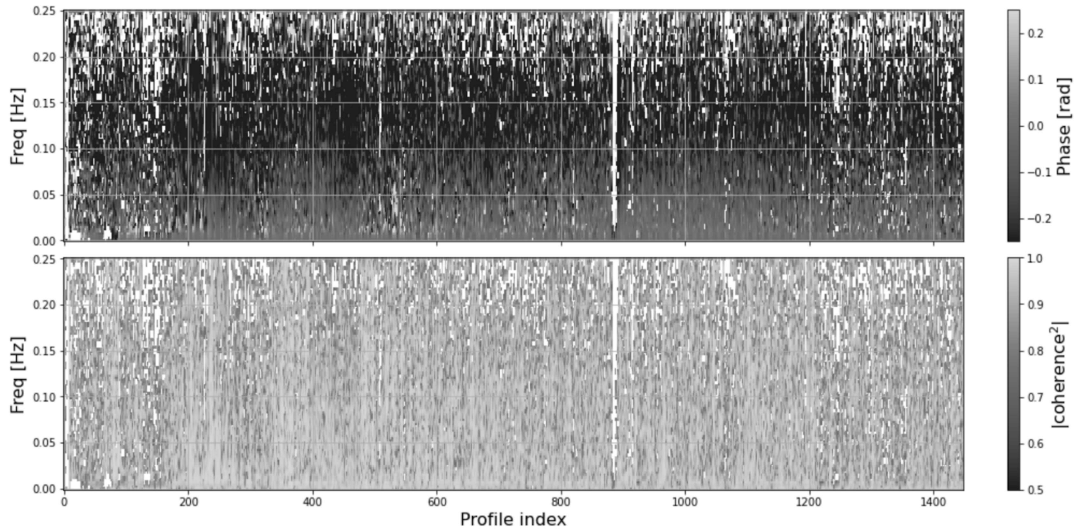

Figure 2.2: Spectrogram for temperature and conductivity cross-spectrum phase (top panel), and squared coherence (bottom panel) as a function of frequency and profile index. Only points exceeding the 95% confidence level and with squared coherence above 75% are shown.

The cutoff density used for this mission was 1023 kg/m^3. This mission had high coherence for most measurements from most profiles, and a similar relationship between cross-spectrum phase and frequency for most profiles (Figure 2.2). The cutoff frequency for the linear fit was chosen to be $f = 0.15$ Hz (Figure 2.3). The resulting value for the time constant $\tau_C = -0.32$ s is well within the required ±0.5 s of the value specified for Sea-Bird GPCTDs, $\tau_C = -0.1$ s. We can thus confirm the Janzen and Creed (2011) constants for the alignment correction are appropriate for use with this GPCTD.

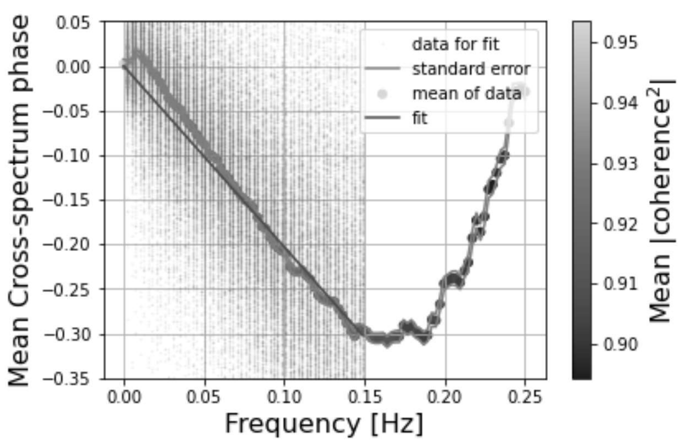

Figure 2.3: Cross-spectrum phase of temperature and conductivity plotted as small gray points, with the mean plotted as large dots coloured by mean squared-coherence. The gray envelope shows the standard error for the mean over all profiles. The red line shows the function fit to points for which the frequency is less than the chosen cutoff and coherence is above 75% and the 95% confidence level is exceeded.

Overall, we do not see large salinity spikes for this mission. However, applying the alignment correction produces a small but quantifiable reduction in small-scale unphysical variations in the salinity field in the region of strong temperature and conductivity gradients (Figure 2.4).

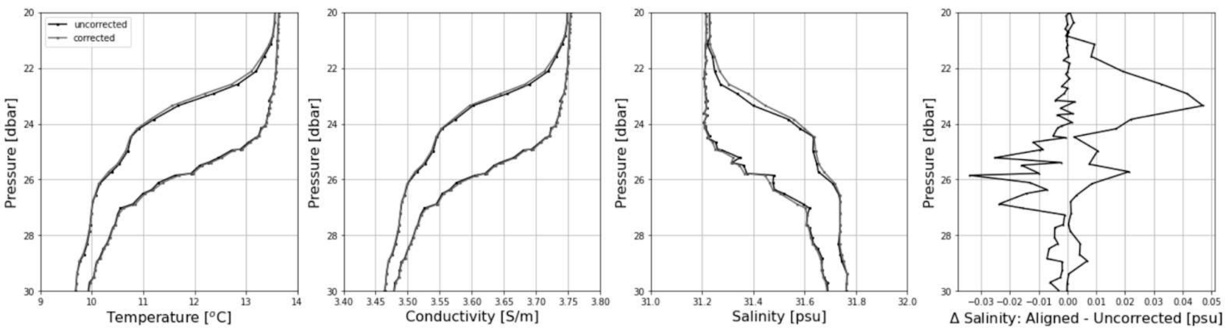

Figure 2.4: Upcast and downcast for temperature, conductivity, and salinity profiles, for part of the upper water column. Black indicates uncorrected fields, and blue indicates fields for which the sensor alignment correction has been applied. Final panel shows the difference between the aligned and uncorrected salinity.

2.3 Identification and removal of questionable salinity profiles

2.3.1 Description of the correction to be applied

We identify any salinity profiles that are obviously unphysical and set all values within those profiles to NaN. Numerous unphysical values of salinity within a single profile are typically caused by something (usually biology) getting caught in the conductivity cell, causing large jumps in the conductivity values recorded. We identify such profiles using a simple criterion applied to the salinity data. The salinity data is first binned by temperature, with bin sizes based on the time series mean temperature profile. This creates temperature bins that are larger in the rapidly varying thermocline and smaller deeper in the water column where temperature varies more slowly with depth.

The criterion applied to the salinity data first temporarily flags any data points that are more than a specified number of standard deviations (typically 4) away from the overall mean for the salinity time series within a given temperature bin, then recomputes the mean and standard deviation, excluding the temporarily flagged values. Salinity values that still differ from the mean by more than the specified number of standard deviations are flagged as ‘bad’. Finally, any profile where more than a chosen percentage (typically 10%) of the salinity values have been flagged as ‘bad’ using this criterion has all values within that profile set to NaN.

This procedure is similar to the ‘Wild Edit’ step available in the SeaBird CTD processing software, but differs in that a full profile is identified as ‘bad’ rather than individual values. The number of temperature bins used may be adjusted, as may the number of standard deviations and the percentage of ‘bad’ values required to identify a profile as ‘bad’.

2.3.2 Application of the correction to mission dfo-mike579-20210704

This mission on the Calvert Line began and ended in Hakai Pass, so that the glider was deployed and recovered close to shore in a coastal region where it is not uncommon to have freshwater pooled in a thin layer at the surface. In order to avoid excluding these low salinity values, the first and last 50 profiles were excluded from the time series to which the correction was applied. We use the standard correction values, with 10 temperature bins, 4 standard deviations, and at least 10% of the profile flagged as 'bad' required to identify it as 'bad'.

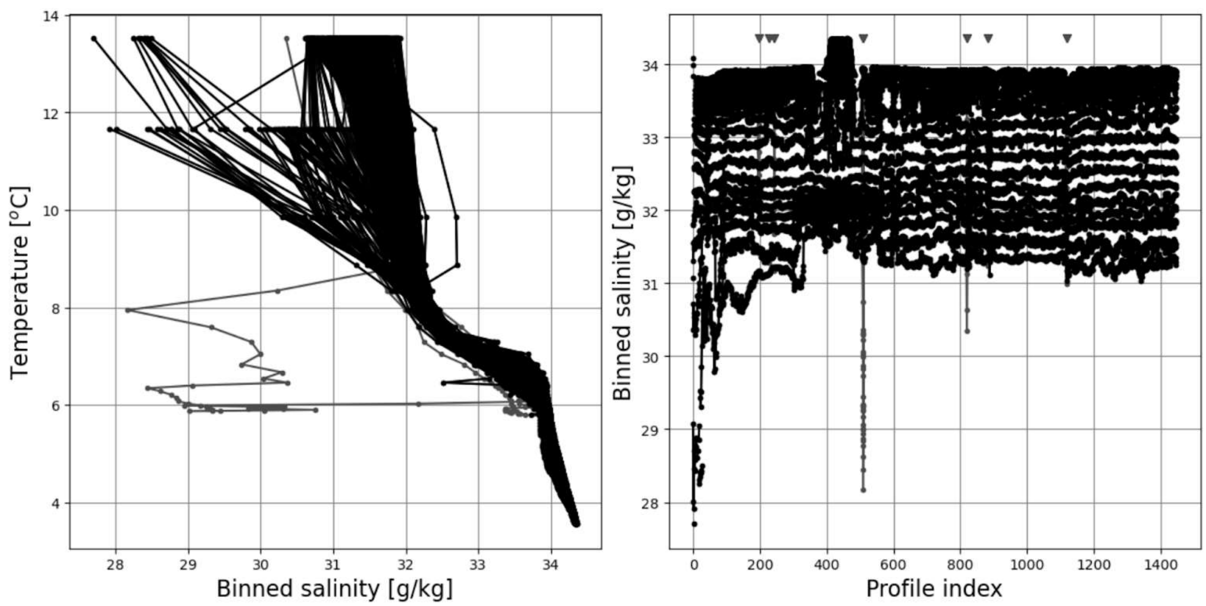

Figure 2.5: Binned salinity plotted as a function of temperature (left) and vs. profile index (right), with the salinity profiles identified as bad due to questionable values and set to NaN shown in red and indicated by the red arrows at the top of the panel on the right.

In Figure 2.5, we see that there are some clearly questionable salinity values. Ultimately, only 8 profiles (numbers 196, 226, 243, 509, 510, 820, 884, and 1119) are identified as 'bad' based on the application of this correction. Setting full profiles to NaN rather than individual values of salinity prevents the loss of measurements associated with real physical phenomena such as upwelling that nonetheless result in significant departures from the mean salinity.

2.4 Thermal lag correction

2.4.1 Description of the correction to be applied

The thermal lag effect is caused by the thermal inertia of the conductivity cell affecting the temperature of the water as it passes through the cell. To estimate the temperature inside the conductivity cell, a recursive filter is applied to the temperature field with parameters α (the amplitude of the error) and τ (the time constant). The values of α and τ that minimize the area between subsequent pairs of profiles, which correspond to the glider diving and then climbing back to the surface, are determined using a brute force minimization scheme. The quantity to be minimized is the root-mean square difference (RMSD), which is calculated as the square root of the sum of the squared areas between pairs of salinity profiles binned by temperature, normalized by the number of pairs of profiles.

The thermal lag correction used here is based on the method proposed by Garau et al., 2011a in “Thermal Lag Correction on Slocum CTD Glider Data”, J. of Atmos. And Oceanic Tech and used in the SOCIB glider data processing toolbox. The Garau et al., 2011a method also minimizes the area between subsequent pairs of temperature-salinity profiles. The Garau et al., 2011a correction was designed for unpumped CTDs, for which the flow rate is not constant. Our Slocum gliders are equipped with Sea-Bird GPCTDs, which are pumped with a constant flow rate. We thus expect the thermal lag to be approximately constant over the full mission.

Garau et al., 2011a states “Morison et al., 1994 showed that there is a relation between the correction parameters α and τ and the flow speed through the conductivity cell. In the case of pumped CTDs, the flow speed is either known or, at worst, can be estimated by observing the misalignment between the sensors’ signals. The flow speed is then assumed to be constant and the correction parameters α and τ are also constant.” As such, it is sufficient to find a single value of α and τ for the entire mission when using a pumped Sea-Bird GPCTD.

Janzen and Creed (2011) determined a cell thermal mass correction for the GPCTD, which they applied to the conductivity data after the sensor alignment correction. For the thermal lag correction, they state, "The applied cell thermal mass corrections were derived using data from prototype glider CTD measurements sampling at 0.5 Hz with a pumped flow rate of 10 ml/s."
The following constants are used for their GPCTD thermal lag correction:

- Amplitude (α) of 0.06
- Time lag (τ) of 10 s

Our procedure for estimating the thermal lag correction is as follows:

1. For each mission, directly estimate the parameters α and τ for the specific GPCTD sensor
2. Confirm that the estimated value of τ is within ±10 s of the Janzen and Creed (2011) value of 10 s
    a. If a given sensor has a directly estimated value of τ that is outside this range, investigate
3. Quantify the improvement for both the Janzen and Creed (2011) parameters and the directly estimated parameters
4. Apply the thermal lag correction with the parameters that result in the greatest improvement, and do not result in an over-correction

Note that the thermal lag correction parameters are more likely to vary between individual GPCTD sensors than are the sensor alignment correction parameters.

To determine the thermal lag correction, a recursive filter on the temperature field is used, where the correction to the temperature, dTn, is calculated from:

$$dTn(n) \;=\; -\,b\,dTn(n-1) \;+\; a\,[T(n) - T(n-1)]$$

here n is the profile index, T is the temperature profile, and a and b are coefficients calculated from:

$$a = 4f_n \alpha\tau/(1 + 4f_n\tau)$$
$$b = 1 - (2a/\alpha)$$

with f_n being the Nyquist frequency (0.5 Hz for the Sea-Bird GPCTD nominally sampling at 1 Hz), α being the amplitude of the error, and τ being the time constant. The goal of the minimization scheme is to determine the values of α and τ that minimize the area between all individual pairs of profiles from a given mission, and then subtract the associated correction to the temperature from the measured temperature before recalculating salinity.

In each iteration of the minimization routine, a polygon is built from two T-S curves, one upcast CTD profile and one downcast CTD profile. The area of this polygon is computed, and then the process is repeated for all pairs of profiles in the time series. The root-mean squared difference (RMSD) is then determined, and the process is repeated for all iterations, varying α and τ within chosen ranges.

The main hypothesis inherent to this correction is that a downcast and subsequent upcast correspond to the same water mass. This assumes a low horizontal advection and small-scale features whose spatiotemporal variability is lower than the resolution provided by a single glider dive. Prior to the minimization, the area between subsequent downcasts is calculated for all profiles, and those for which the area is greater than one standard deviation from the mean are excluded from the minimization routine. This preliminary step ensures that we do not include pairs of profiles where the glider passed through an eddy, a front, or other small-scale features.

Furthermore, a cutoff is imposed on the area between pairs of profiles that are included in the subset used for the minimization scheme. Any pair of profiles whose area is more than 3 standard deviations away from the mean will be excluded from the determination of the RMSD. This ensures that a small number of anomalous profiles do not bias the results.

2.4.2 Application of the correction to mission dfo-mike579-20210704

This mission on the Calvert Line occurred in a highly energetic environment, so near-surface differences between a downcast and the subsequent upcast are likely to be caused by spatiotemporal variability. The glider dove to a maximum depth of 1000 m, taking several hours to resurface. During this time, near-surface water masses or small-scale features present were likely to change. As such, we exclude near-surface segments of each profile for which the density was <1023 kg/m^3 from the minimization routine. We also exclude the first and last 50 profiles, which were primarily collected in the highly energetic environment of Hakai Pass. The effect of isopycnal heaving on the profiles is not a concern, since the minimization is conducted in temperature-salinity space. We use a subset of the remaining data, consisting of 100 pairs of profiles equally spaced in time, to determine the correction to be applied.

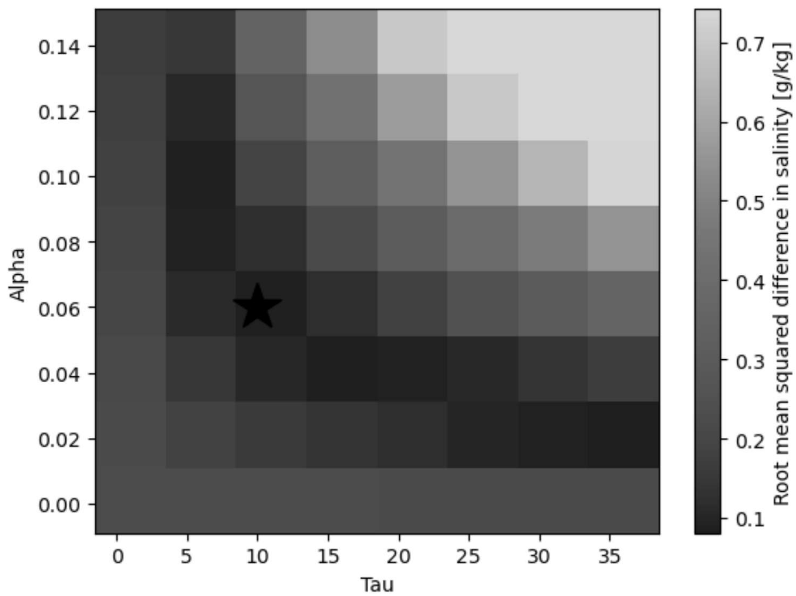

Figure 2.6: Values of the root-mean squared difference (RMSD) for each value of α and τ used during the grid search to minimize the area between pairs of profile. Note that high values of the RMSD have been saturated (yellow) to better show the minimum (dark blue). The black star indicates the Janzen and Creed (2011) values.

The values of α and τ are estimated using a simple grid search over a range of possible values. Values of $\alpha \times 1000$ between 1 and 150 were tested in increments of 20, and values of τ between 1s and 35s were tested in increments of 5 s (Figure 2.6). It was found that $\alpha = 0.02$ and $\tau = 35$ s best minimized the RMSD, with RMSD = 0.081ºC, a decrease of 0.173ºC from the original RMSD = 0.254ºC. We see in Figure 2.6 that there is a fairly wide band of 'good' values where the RMSD is low, with the Janzen and Creed (2011) values well within this band. We additionally confirm that this band of 'good' values is the same for profiles taken from 1) the early part of the mission, using profiles from the first quarter of the mission, and 2) the latter part of the mission, using profiles from the last quarter of the mission.

Since the directly determined parameters fall in the same band of low RMSD values as the Janzen and Creed (2011) parameters, i.e. they result in only marginally lower RMSD overall compared to the Janzen and Creed (2011) parameters, quantifying the improvement for our directly determined values is deemed unnecessary. The Janzen and Creed (2011) parameters are used for the final thermal lag correction.

The thermal lag correction results in a clear qualitative improvement in the salinity field, with a reduction in the up-down asymmetry throughout the water column (Figure 2.7). The quantitative improvement that results from applying the thermal lag correction is shown in Figure 2.8, which demonstrates a decrease in the area between most individual pairs of profiles after the correction was applied, for all profiles pairs collected during the mission. The improvement is most significant for full-depth profiles to 1000 m, with similar improvement seen on the outbound and inbound legs of the mission. The magnitude of the difference between the corrected and uncorrected fields is over 0.05 psu in the thermocline for the full-depth profiles, which is significant given that the typical range of salinities measured in this region spans ~15 psu.

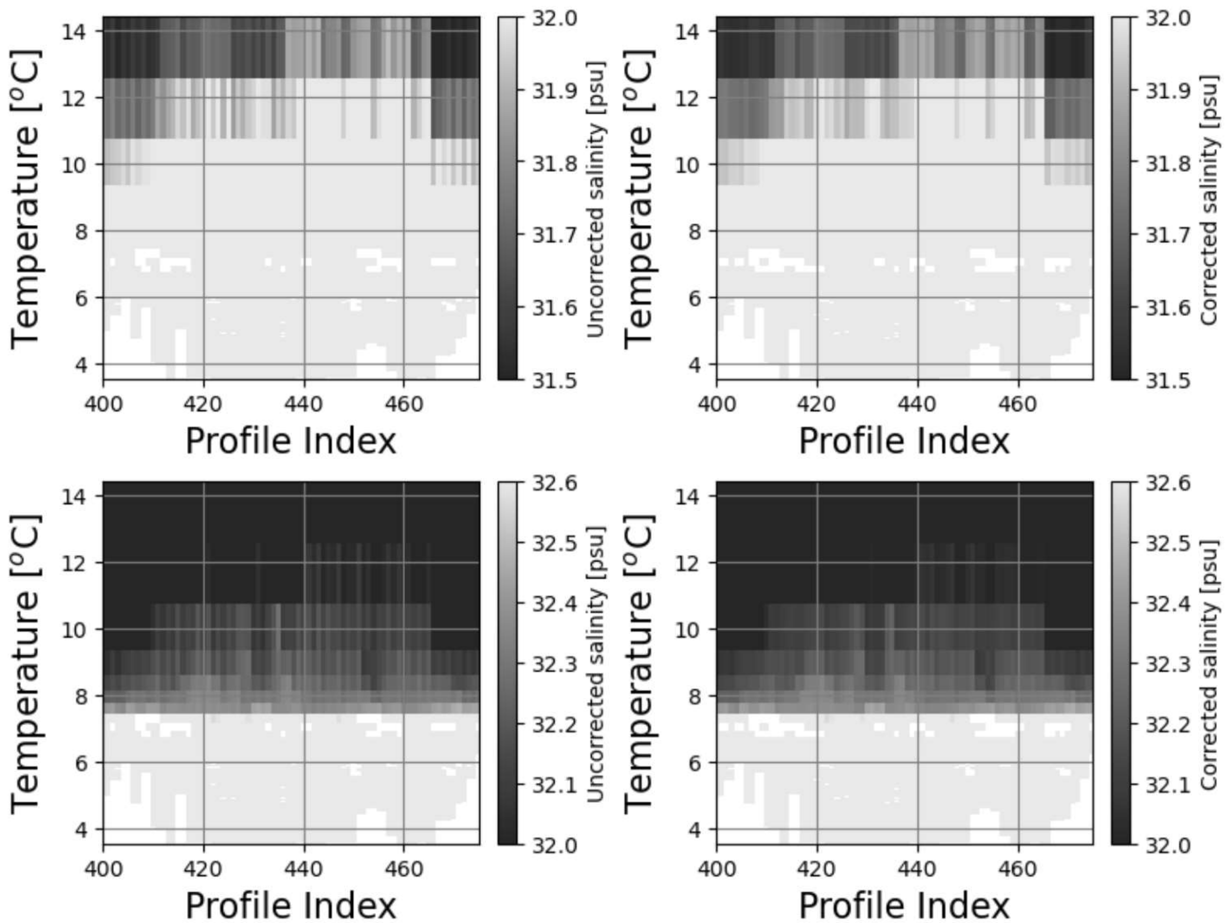

Figure 2.7: Salinity binned with temperature and plotted vs. temperature and profile index. Profiles deeper than 1000m are shown. The uncorrected salinity is shown in the left panels and the corrected salinity using the Janzen and Creed (2011) parameters is shown in the right panels. Note that the colorbar limits for the bottom two panels have been adjusted to show variations deeper in the water column.

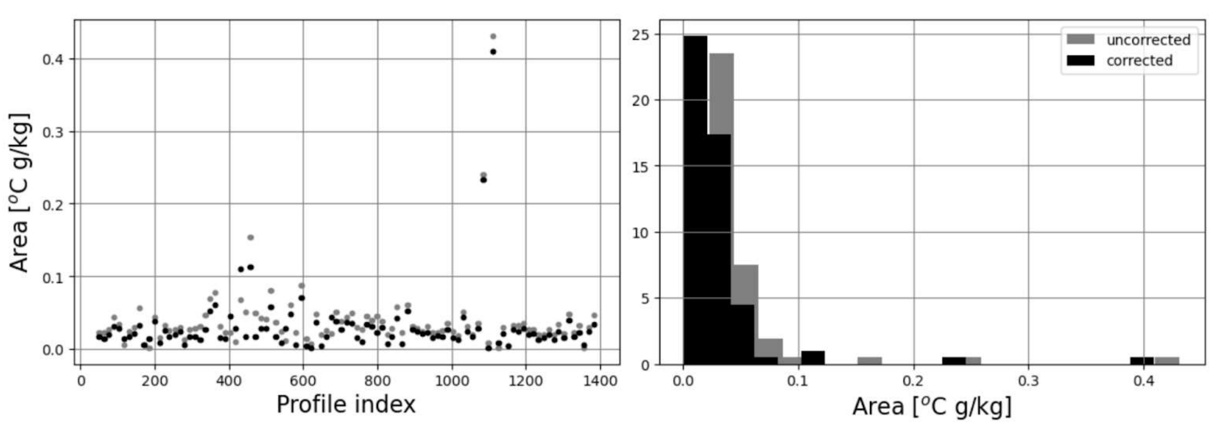

Figure 2.8: Area between pairs of salinity profiles, calculated in temperature-salinity space, plotted vs. profile index number (left) and as a histogram (right), for the uncorrected salinity field (grey) and the corrected salinity field (black).

Finally, the effect of applying the thermal lag correction is shown in Figure 2.9 in temperature-salinity space, for a subset of profiles. In the data without the thermal lag correction applied, there is a small offset between the downcasts and upcasts, particularly in the thermocline. The upcasts are less saline throughout much of the water column, though the offset is significantly smaller at densities greater than about 1025.0 kg/m^3. After the correction is applied, this offset is visibly reduced, particularly within the thermocline, and there is no evidence of an over-correction.

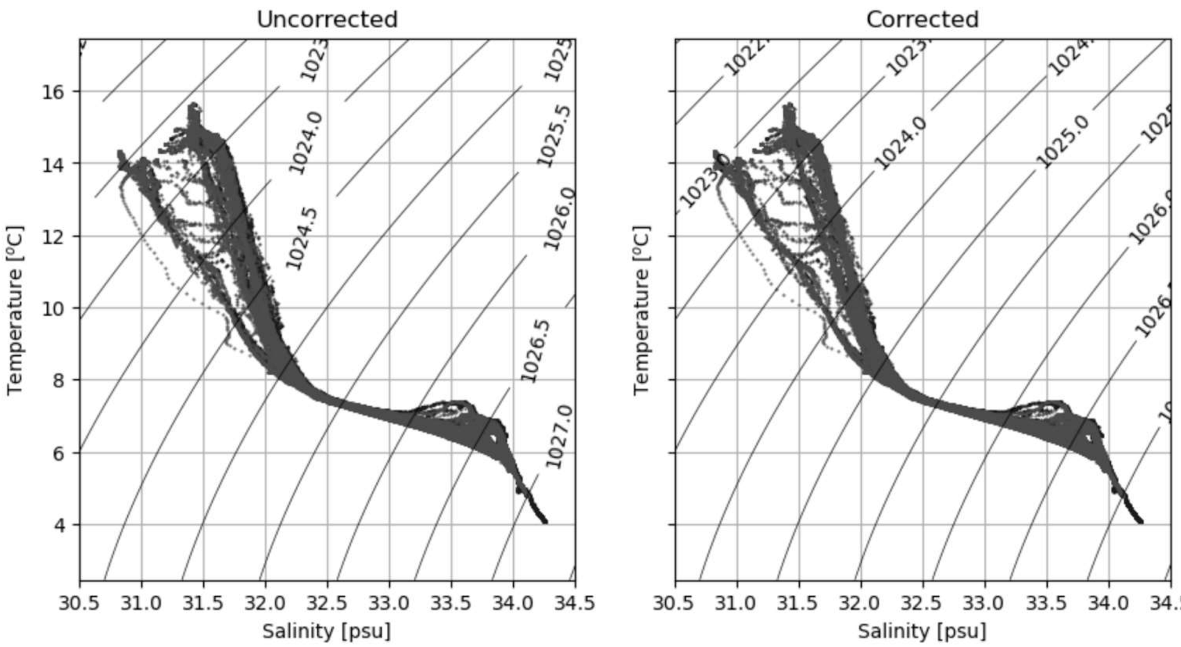

Figure 2.9: Temperature-salinity diagrams for a small number of profiles, showing the difference between upcasts (red) and downcasts (blue), for the data without the thermal lag correction applied (left panel) and the data with the thermal lag correction applied (right panel).

3 Summary of corrections applied

Ocean gliders in the C-PROOF fleet have the following corrections applied to their Sea-Bird GPCTD conductivity-temperature-depth sensors, in order:

- Identification of anomalous conductivity values
- Sensor alignment correction
- Identification of questionable salinity profiles
- Thermal lag correction

The corrected temperature and salinity fields are shown in Figure 3.1.

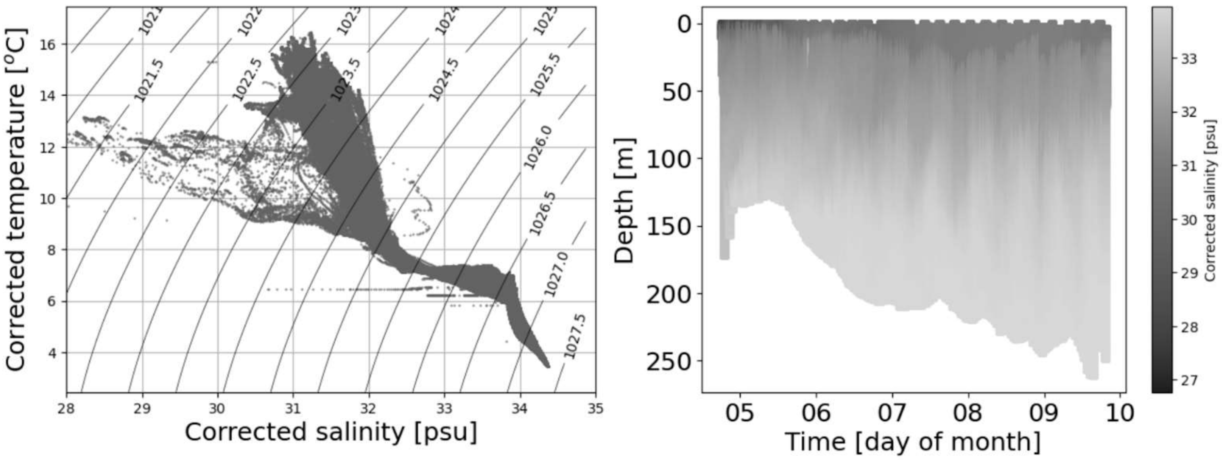

Figure 3.1: The corrected temperature and salinity fields shown in a T-S diagram with density contours (left) and the corrected salinity field plotted vs. depth and time (right) for mission dfo-mike579-20210704.

References

1. Ferrari, R., and Rudnick, D. L. Thermohaline variability in the upper ocean, J. Geophys. Res., 105(C7), 16857-16883, 2000.

2. Garau, B., Ruiz, S., Zhang, W. G., Pascual, A., Heslop, E., Kerfoot, J., & Tintoré, J. Thermal Lag Correction on Slocum CTD Glider Data, J. Atmos. Oceanic Technol., 28(9), 1065-1071, 2011.

3. Janzen, C. D., and Creed, E. L. Physical oceanographic data from Seaglider trials in stratified coastal waters using a new pumped payload CTD, OCEANS'11 MTS/IEEE KONA, Waikoloa, HI, USA, 1-7, 2011.

4. Morison, J., Andersen, R., Larson, N., D'Asaro, E., & Boyd, T. The correction for thermal-lag effects in Sea-Bird CTD data, J. Atmos. Oceanic Technol., 11, 1151-1164, 1994.

5. Sea-Bird Seasoft V2:SBE Data Processing - CTD Data Processing & Plotting Software for Windows, Sea-Bird Scientific, software manual revision 7.26.8, 2017.

6. Sea-Bird User Manual - GPCTD Glider Payload CTD (optional DO) - Conductivity, Temperature, and Pressure (optional DO) Sensor with RS-232 Interface, Sea-Bird Scientific, manual version 008, 2021.