Determining Seaglider® Velocities Automatically

James Bennett, Fritz Stahr, and Charlie Eriksen
School of Oceanography, University of Washington
August, 2019

Introduction

Determination of underwater glider speed along its glide path is critical to recovering important oceanographic observations, including vertical water velocities and depth-averaged currents (DAC), as well as for properly correcting passive flow-dependent sensors, including thermal-inertia effects in unpumped conductivity-temperature (CT) instruments.

Direct measurement of vehicle speed along its glide path is typically unavailable[1]. However, pressure is continually measured providing an indication of a vehicle's vertical velocity (*w*). Given a model of steady flight for a vehicle in relatively still water, multi-variable regressions against the measured *w* can determine a set of vehicle-specific flight parameters (e.g., coefficients of lift and drag, etc.) that minimize discrepancies between predicted and measured *w* for various dives. These parameters are then assumed to characterize steady vehicle flight and permit estimates of horizontal and vertical velocities; any residuals presumably indicate oceanographic processes operating on the vehicle, such as internal waves, turbulence, and depth-averaged currents.

Solving the required regressions has been left to the pilot or oceanographer, which, if attempted[2] using poorly-documented versions of the `regress_vbd` MATLAB script[3], requires careful selection and preparation the profile data and knowledgeable interpretation of the regression results. If the results of an improperly-performed regression are applied to the glider during operation (via parameters `$HD_A`, `$HD_B`, `$HD_C`, etc.), the vehicle can incorrectly calculate desired pitch and buoyancy values, impacting navigation and flight.

Further, even when the regressions were performed well, earlier basestation[4] processing assumed that a single characterization of the flight model sufficed for the

---

[1] But see Merckelbach, et al., 2019, Todd, et al., 2017, and recent experimental results with velocimeters onboard Deepgliders, discussed below.
[2] Many pilots rely on a set of default flight model parameters for the vehicle type. However, as we will see, vehicle construction differences and sensor suite changes make trusting even our improved default parameter values problematic.
[3] Including code derived from it, such as Bastien Queste's Seaglider Toolbox.
[4] The code base that processes raw data files from the glider during and after a deployment into a quality-controlled format for further oceanographic analysis.

entire deployment. However, this approach is unable to account for changes to vehicle shape or buoyancy due to biofouling or damage during the deployment.

To address these and other issues we developed a completely automated flight model system (herein 'FMS') that solves the flight model regressions regularly and consistently recovers well-constrained flight model parameters *for each dive* of a deployment. It is designed to analyze both on-going and completed deployments, providing, in the former case, warnings and suggestions to the pilot about flight performance and parameter settings that could improve glider's navigational performance at sea.

Frajka-Williams, et al., 2011 (herein 'FW2011') outlined a similar regression procedure for Seagliders and demonstrated its ability to recover accurate vertical water velocities when applied to two deployments in the Labrador Sea in 2004 and 2005. Here we automate and extend their procedure and observations, generalizing it to apply to all Seaglider®, Deepglider®, and Oculus Coastal® glider types[5].

After reviewing the Seaglider flight model, we outline our automated regression procedure and discuss its operation on a number of interesting deployments, including, for comparison, the Labrador Sea deployments mentioned above. Various operational details and considerations are reserved for an appendix.

The Seaglider Flight Model

Here, we rehearse FW2011's description of the flight model (Eriksen, et al.,2001) used to characterize Seaglider-class vehicles in steady flight. We maintain and extend their notational conventions for easy comparison.

Consider a glider flying at a given glide angle θ as shown in Figure 1.

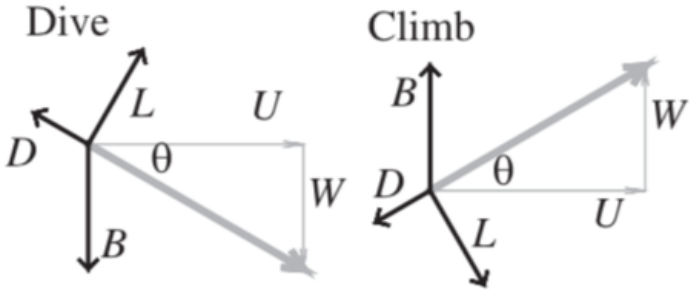

[5] These are commercially available from Kongsberg Underwater Technology as models Seaglider M1, Seaglider M6, and Seaglider C2.

Figure 1: Diagram of the force balance on a Seaglider during a dive (left) and during a climb (right). Forces are lift (L), drag (D), buoyancy (B). Glider velocity is U and W in the x and z directions, and θ is the glide angle. Reproduced from FW2011.

The flight model assumes, in steady flight, that the vertical lift L, drag D, and buoyancy B forces are in balance:

$$B + L\cos\theta - D\sin\theta = 0 \tag{1}$$

The following equations are assumed to describe the lift L, drag D, and buoyancy B, forces.

$$L = ql^2a\alpha \tag{2}$$

$$D = ql^2(bq^s + c\alpha^2) \tag{3}$$

$$B = g(\rho V(t,p,T) - M) \tag{4}$$

where l is drag-scale factor[6], a the lift coefficient, α the attack angle, b the drag coefficient, c the induced drag coefficient, g the gravitational acceleration, q the dynamic pressure, s the impact of the hull shape on q, ρ the measured in-situ water density, M the measured glider mass, and $V(t,p,T)$ the volume of the glider, which depends on t time, p measured pressure, and T measured water temperature.

The measured vehicle pitch ϕ is the sum of the glide angle θ and the vehicle's angle of attack α: ϕ = α + θ. The dynamic pressure is equal to $q = \rho_0(U^2 + W^2)/2$ where ρ_0 is a reference density, here fixed for all vehicle types at a nominal 1027.5 kg/m^3, and U and W are the horizontal and upward-positive vertical vehicle speeds relative to water motion. The impact of hull shape on drag, s, has been estimated from previous work: s=−1/4 for the tapered low-drag Seaglider shape (Hubbard, 1980), s=0 for the Deepglider and Oculus standard right-cylinder hulls with appended fairings.

The buoyancy force B results from the difference between the mass of the glider M and the seawater displaced by the glider volume V. The total volume V varies over the course of a dive due to the glider's hydraulically-pumped variable buoyancy device (VBD), the expansion of the hydraulic oil with temperature, the compressibility of the hull and sensors with pressure, and the expansion of the hull and sensors with temperature. In addition, for vehicles employing compressible fluids (CF) for improved buoyancy performance, the effective volume changes due to density changes of the CF with pressure and temperature.

---

[6] The proportionality factor between hydrodynamic forces and dynamic pressure has dimensions of area. Following conventional practice, we chose a nominal hull length (1.8m) l squared for all vehicles, permitting comparison of relative lift and drag coefficients found between vehicles. Further this length value is a fixed constant in the glider flight code shared between these vehicles.

$$V(t,p,T) = V_{hull}(t)e^{-(\kappa p - \tau(T-T_0))} + V_{CF}(p,T) \quad (5)$$

$$V_{hull}(t) = (V_0 - \delta V_0) + V_{VBD}(t) - M_{CF}/\rho_{CF}(0,T_0) \quad (6)$$

$$V_{CF}(p,T) = M_{CF}/\rho_{CF}(p,T) \quad (7)$$

where $V_{VBD}(t)$ ($V_c(t)$ in FW2011) is the volume of the VBD system around a nominal reference point, V_0 is a (computed) total reference volume, T_0 is an arbitrary reference temperature (15°C for all vehicle types), κ (α_g in FW2011) is the overall compressibility of the combined hull, any syntactic foam or foam-filled fairing elements, and sensors for the vehicle, τ (γ_g in FW2011) the volumetric thermal expansion, M_{CF} is the mass of any compressible fluid and $\rho_{CF}(p,T)$ is its density at pressure p and temperature T. For nearly all Seaglider (and Oculus) deployments, M_{CF} is zero. The δV_0 bias term captures unmodeled variation in volume (strictly, density[7]) due to, for example, oil shrinkage due to decreasing temperature, water uptake by the fiberglass fairing elements, biofouling, loss of CF, etc. The convention of subtracting δV_0 from V_0 is arbitrary and historical; increasing δV_0 implies a decrease in overall vehicle volume and hence an increase in its density given constant M.

The system varies the flight parameters to find a set that best characterizes the flight regime of the vehicle. Given a candidate set of flight parameters and an estimate of buoyancy B, Equations 2 and 3 are solved iteratively for q and α that satisfy Equation 1. Thus, buoyancy serves to constrain the values of q and α and hence vertical and horizontal speeds. From q and α, $U^2 + W^2$ and θ are determined, from which we can estimate the modeled vertical velocity w_{stdy} as

$$w_{stdy} = \sqrt[2]{(U^2 + W^2)} \sin\theta \quad (8)$$

During the regressions, the modeled w_{stdy} is compared with $w = dz/dt$, the measured glider vertical velocity. Glider vertical position z, which is measured positive downward, is computed from measured pressure p. The net vertical velocity of the water w_w is estimated by:

$$w_w = w - w_{stdy} \quad (9)$$

Overall the system attempts to find a set of flight parameters that minimizes the root-mean-squared (RMS) differences between the measured and predicted w:

[7] In general, there are no independent measurements of long-term un-commanded changes in vehicle density. In our model these changes are reflected as volume changes against an assumed constant mass. Alternatively, these changes could be reflected as changes in mass against a constant volume; this is the approach taken by Rudnick et al., 2013.

$$w_{rms} = \sqrt[2]{\overline{w_w{}^2}} \tag{10}$$

over the points in steady flight, assuming zero net vertical transport in the water column. This model handles neither accelerations during VBD operation nor side-slip due to chronic or commanded banked flight, unlike Merckelbach et al.'s 2019 dynamic flight model; those effects are discussed more below.

The FMS Procedure

Given pre-deployment measurements of the overall vehicle mass M and any compressee mass M_{CF}, the goal of the FMS is to recover—via various regressions and using per-dive measurements of V_c, pitch ϕ, temperature T, pressure p, and in-situ seawater density ρ—accurate values for the remaining flight parameters for each dive as the deployment progresses, and hence recover the per-dive glider speeds and glide angles.

The flight equations describe the vehicle's steady flight in still water. To accurately estimate the flight parameters, FMS employs heuristics to select dive data points that reflect steady flight. Measurements are discarded when the VBD engine is accelerating the vehicle during dive, apogee, and climb. While FW2011 observed that some rolls (on SG014 in the Labrador Sea, their figure 6) could momentarily accelerate the vehicle, this was likely due to pitch/roll mechanism coupling in that version of the vehicle. Analyzing rolls over many deployments we found this effect to be rare, small, and (even in the deployment they analyzed) negligible. Thus, data during rolls (and small pitch changes under auto-pitch adjust operation) are retained, which permits, for example, compass-calibration dives employing constant roll to be used in the flight model regressions.

The seawater temperature and in-situ density used during the regressions are estimated from the uncorrected values returned from the CT instrument, before any speed-dependent thermal-inertia adjustments are performed; the typical impact of these adjustments for the purposes of estimating the flight parameters was also found to be negligible. However, any conductivity anomalies (bubbles, etc.), electrical spikes in temperature, etc. are removed. Other quality-control heuristics ensure the pressure sensor, compass (for pitch), and the CT system are operating properly before using a dive's data. Finally, the system selects points where the measured w is changing slowly, indicating relatively quiet water.

As dives are received, the flight regressions proceed in two steps. The first step determines the buoyancy force B for each dive by estimating the vehicle's V_0 then per-dive values for δV_0 and κ. The second step estimates the current flight regime itself, lift a and drag b in particular, using subsets of adjacent dives selected at regular intervals. These steps are described in detail below using the relatively

uneventful 'normal' deployment of SG144 off Ocean Station PAPA in June 2009 (Pelland et al., 2016). This eight-month deployment maintained a bow-tie pattern centered on Ocean Station PAPA until dive 660, when she flew 1300km to the Washington coast for recovery.

The two-step regression structure of FMS follows the procedure outlined by FW2011 and represents a simplified version the single, full, multi-variable regression originally implemented in MATLAB. Accurate initial determination of B has the largest impact on vertical velocity estimates compared with lift and drag estimation. Following FW2011, various sensitivity analyses indicates that variance in several parameters have negligible impact on speed estimates and can thus be estimated once and fixed for each vehicle type: τ is dominated by the hull material, and c and s are both determined by hull shape. The induced drag parameter c was empirically estimated for each vehicle type by investigating near-stall and flat-spin behavior of some dives; however, changing c even by an order of magnitude has very little impact on the a/b choice found by FMS[8].

Determining Buoyancy Forcing: Estimating V_0, δV_0 and κ

Accurate determination of buoyancy forcing B for each dive requires estimating the buoyancy parameters V_0, δV_0, and κ in Equation 4. This is done first by estimating a likely value of V_0 from the vehicle's mass and measured seawater density at apogee (discussed below). Then, assuming a current best-guess for lift and drag parameters, FMS regresses δV_0 and then κ in order to minimize dive's w_{rms}. If the best guess for lift and drag change subsequently (see below) δV_0 and κ are automatically re-estimated.

FMS produces summary figures, updated during operation, showing the trends in δV_0 (Figure 2) and κ (Figure 3) over the deployment assuming an estimated V_0.

[8] The original MATLAB version of the regressions often found values for c several orders of magnitude smaller than our default estimates. When these values are placed aboard the glider, the single-precision floating point arithmetic on the TT8 (and RevE) processor incorrectly estimates speeds and often requests positive (nose-up) pitch angles for dives.

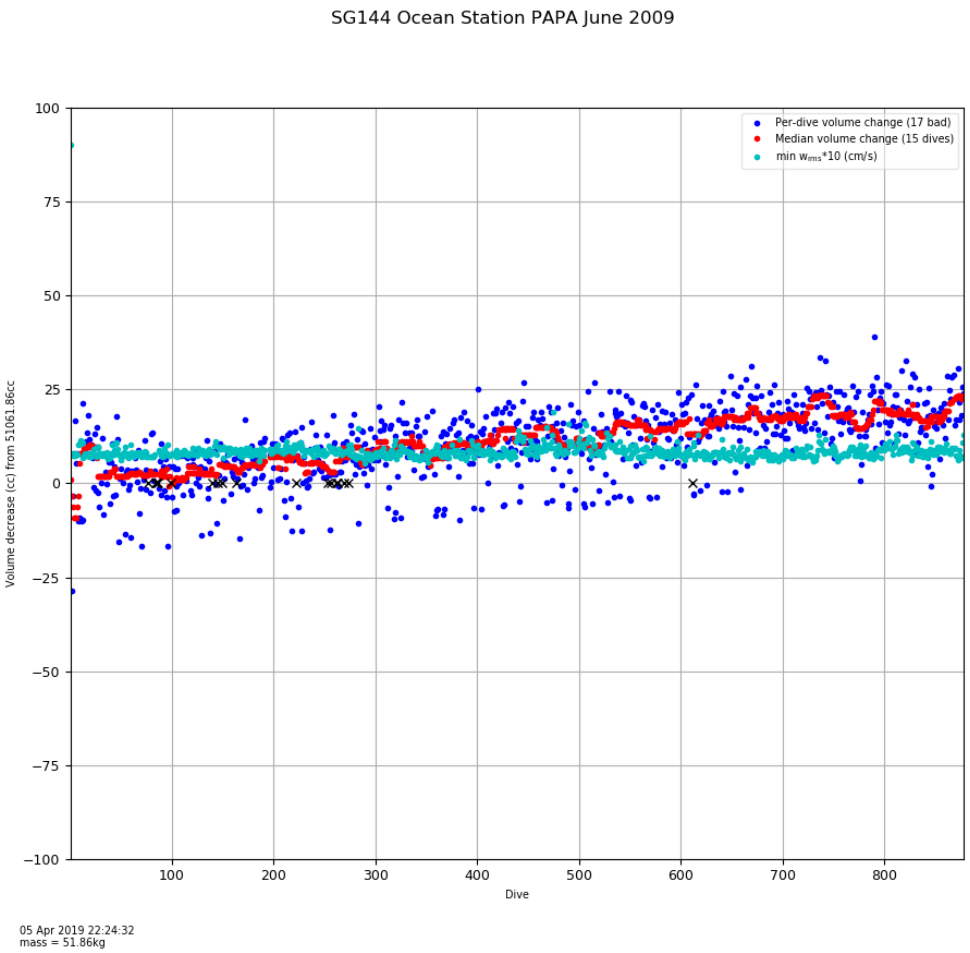

Figure 2: Per-dive solutions of δV_0 for SG144 off Ocean Station PAPA in June, 2009. V_0 was determined to be 51061 cc. Blue points are per-dive estimates. Red points are a 15-dive median filtered trend. Cyan points show scaled per-dive w_{rms} values; these are relatively constant at 0.87 cm/s.

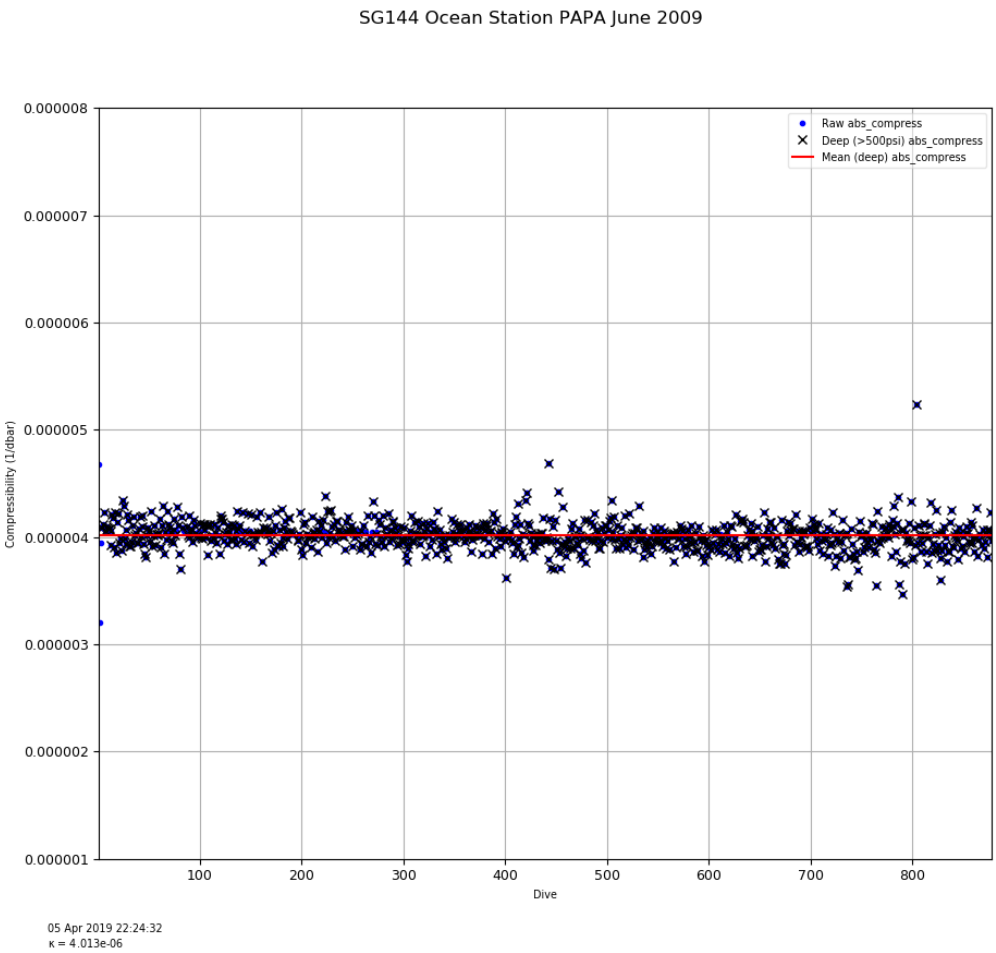

Figure 3: Per-dive solutions of compressibility κ for SG144 off Ocean Station PAPA in June, 2009. Nearly all dives were to 1000m. Mean κ was determined to be $4.013e^{-6}$.

As discussed in FW2011, changes in (V_0 - δV_0) adjust estimated w_{stdy} on dive and climb around a mean value for a given κ and a/b regime; the first regression for δV_0 minimizes that mean difference. Once δV_0 is determined, the κ regression ensures that w_{stdy} estimates in the deep portions of the dive (> 500 psi, where higher pressures have a larger, secondary impact on the volume of the hull) reduce w_{rms} further.

FMS implicitly assumes that the mean w_w during a dive and over the entire deployment is zero. While strong upwelling and downwelling signals in $w_w$are frequently seen, they are transient compared with the dive duration. Typical long-term vertical shear in the ocean is estimated at ~0.5-1.0 m/d (Pelland et al., 2017), or $\sim 1e^{-5}$m/s, significantly below the resolution provided by FMS given typical data noise.

The δV_0 graph for this deployment shows a high-frequency variation around a long-term mean trend of decreasing volume. The high-frequency variation is due to the

regression overfitting internal-wave signals in w that do not cancel over the course of a dive as well as noise in both pressure and pitch sensors. This can be seen in Figure 4, which shows a roughly ± 12 cc variation produced around each per-dive δV_0 for w_{rms} solutions within 0.04 cm/s of the per-dive minimum.

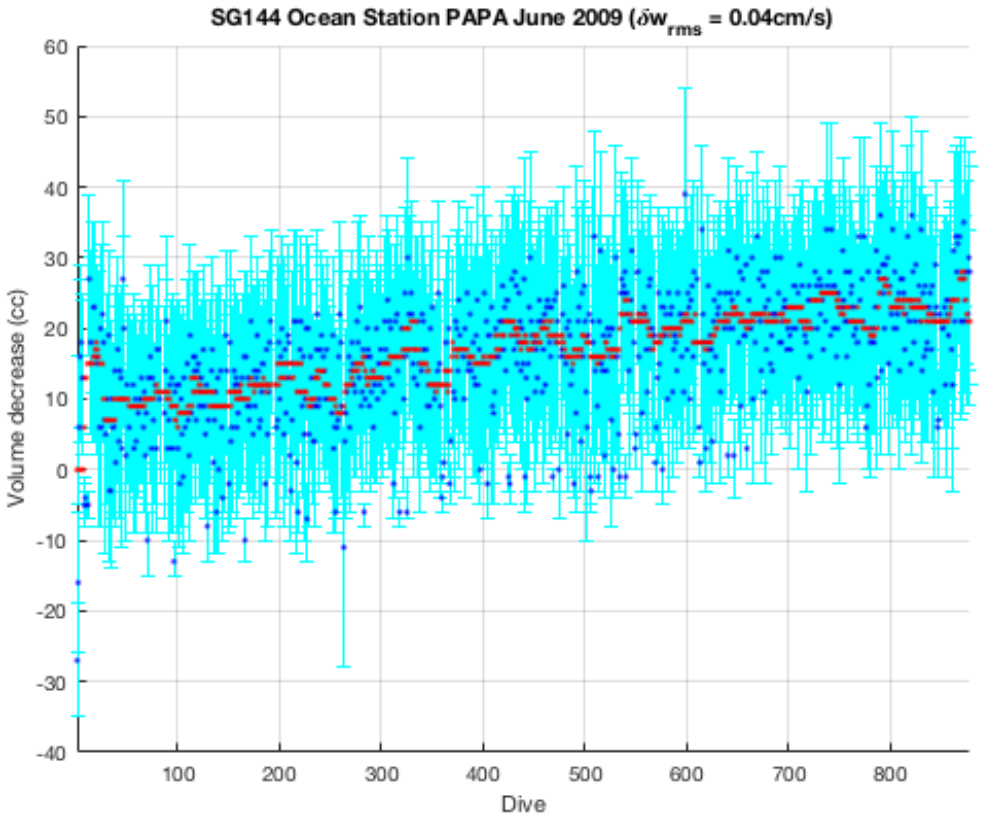

Figure 4: Variation in volume (cyan bars) with respect to the minimum per-dive solutions (blue dots) of δV_0 for a small tolerance (0.04 cm/s) around the per-dive minimum w_{rms} for SG144 off Ocean Station PAPA in June, 2009. Red points are a 15-dive median filtered trend of the minimum δV_0 values.

We conjecture that the long-term mean decrease in volume is, in the case, due to slow water intrusion into the fiberglass and syntactic foam outer fairing of the glider; it is seen on several other deployments (not shown). In spite of these changes the mean w_{rms} for the minimum solutions was 0.87 cm/s, consistent with the relatively quiet deep waters off Ocean Station PAPA.

While κ was originally expected to be dominated by the compressibility of the hull material (aluminum or carbon fiber) and could, perhaps, be treated as a constant per vehicle type, in fact, the mean value was observed to change slightly per vehicle

and per deployment. This vehicle variation likely reflects the different, combined compressibility of the hull, sensors, and fairing components[9].

Estimating Vehicle Reference Volume: V_0

The δV_0 regression operates against a prevailing estimate of V_0 for the vehicle. Estimating the overall reference volume V_0 is done in two steps[10]. While processing the first dive, FMS estimates V_0 via

$$V_{o,initial} \approx \frac{M}{\rho_{apogee}}$$

where ρ_{apogee} is the measured in-situ seawater density at apogee when $V_{VBD}(t) = 0$. This initial assumption is often poor: typically, there are bubbles trapped in the fairing that must be dissolved and the VBD system is not always adjusted to reflect true neutral buoyancy at apogee. Nevertheless, with a $V_{0,initial}$ estimated the system is able to then estimate a δV_0 such that ($V_{0,initial}$ - δV_0) reflects an accurate volume, hence density of the vehicle and therefore B. Once the system estimates δV_0 values for several additional dives against $V_{0,initial}$ it then re-estimates a final V_0 that would reduce the mean δV_0 value to zero. The contributing dives are then reprocessed against these modified (V_0 - δV_0) values and all subsequent dives assume this final V_0. In the case of SG144, $V_{0,initial}$ was estimated at 51212 cc with δV_0 values around 150 cc; $V_0$was then adjusted to 51061 cc after four dives. This value compares well with the pilot-supplied value of 51075 cc.

It is important to stress that it is the sum (V_0 - δV_0) that reflects the per-dive volume of the vehicle at neutral density for accurate buoyancy estimation; restating V_0 to reduce initial mean δV_0 values to zero is a convenience when reviewing subsequent δV_0 trends.

Poor estimates of V_0 (for the stated mass M) in `sg_calib_constants.m` have been the single largest source of problems with basestation operation, often leading to incorrectly processed profiles in which dives or climbs or both appear to be lost due to induced apparent 'stalls' arising from bad buoyancy estimates. FMS eliminates these problems. Indeed, the FMS procedure for estimating V_0 does not require an accurate vehicle mass M. As long as the measurement of in-situ seawater density is accurate, V_0 (and its associated per-dive δV_0) will be scaled properly to

[9] The construction of the VBD engine in the Oculus glider precludes accurate estimation of the vehicle's κ (but not δV_0) so this step is skipped and a nominal value is used. The external piston moves under pressure, dominating whatever hull compression occurs; the piston thus appears as a very squishy hull element.

[10] There is a preliminary estimate of V_0 using the mass M and the fixed constant ρ_0, before any measurement of density is available; this is rapidly supplanted by subsequent measurements.

yield accurate buoyancies[11]. However, accurate measurement of M_{CF} is critical since the CF's absolute volume change (second term in Equation 5) is scaled by M_{CF} using the CF's equation of state $\rho_{CF}(p, T)$. Of course, the system assumes the CF's equation of state has been accurately determined over oceanographic temperature and pressure regimes (see Bennett, et al. 2018 for an experimental determination of equations of state for typical CFs).

Determining Lift/Drag Flight Regimes

FMS next combines dive data and buoyancy B estimates from subsets of adjacent dives to estimate a single lift a and drag b pair that minimizes the combined w_{rms} of the dives. To compare solutions between different sets of dives and visualize the range of acceptable a/b pairs, FMS implements this 'regression' by solving for w_{rms} at fixed a/b grid points. The result of a typical single grid search is shown in Figure 5; FMS generates these figures for each grid search it performs.

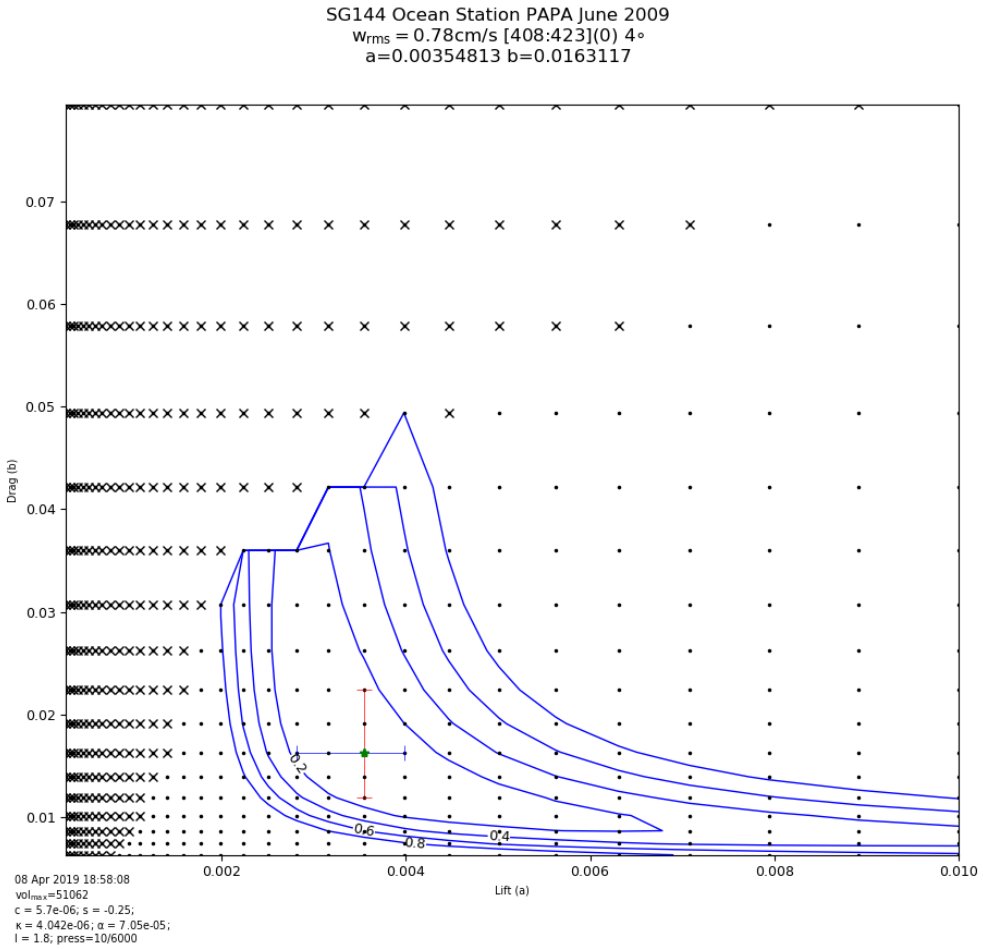

[11] Easily verified by processing the same deployment with different stated masses and ensuring that, while V_0 changes, the regressed flight parameters and speed estimates remained unchanged.

Figure 5: Lift/drag grid solution space of w_{rms} combining dives 408 thru 423 from SG144 off Ocean Station PAPA in June 2009, providing 4° pitch separation. Minimum $w_{rms}$is 0.78 cm/s at the green dot. Contour lines show increases of w_{rms} of 0.2 cm/s. Vertical and horizontal bars indicate w_{rms} solutions within 0.2 cm/s centered on the minimum point. Excessive stalls (and hence no velocity solutions) are indicated by a black x.

The choice of grid points reflects empirical observations of typical solution contours. Overall the minimized w_{rms} solutions resembles a shallow bowl where most of the minimum solutions lie in the lower-left corner of the grid space during typical operation. However, as some vehicles become biofouled their drag coefficients rise (see below). In addition, very low lift, high drag pairs in the upper-left corner of the space often lead to substantially stalled solutions (indicated by black x's).

FMS performs a grid search at regular and frequent intervals, typically every 16 dives[12], permitting detection of major changes in lift and drag coefficients. If the lift and drag values change from the previous best guess from the last grid search, each intervening dive's δV_0 and κ is provisionally re-computed using the new values of a and b; the a/b with the lowest w_{rms} is applied to each dive and the dive is reprocessed if needed using changed flight parameters. In this way all dives are constantly updated to reflect the best composite parameter estimates.

FW2011 observed that the drag coefficient b can be better constrained by combining dives with very different pitches (attack angles). FMS attempts to combine recent dives that maximize the spread of pitches routinely obtained by those dives. However, to minimize the time to compute a grid search, FMS selects a subset of the dives that maximize the pitch spread of those dives and minimize the number of total data points used in the search[13].

Empirically we found that pitch spreads exceeding roughly 7° provide well-constrained values of b; we label these grid solutions 'trusted'. Figure 6 below shows the subsequent grid solution to the solution given in Figure 5 above. The pitch spread in Figure 5 is only 4° and the range of acceptable drag values is large; the pitch spread in Figure 6 is substantially larger at 28° and the values of b are much more constrained. However, the value of b in Figure 6 is nevertheless within a 0.2 cm/s w_{rms} difference contour of Figure 5 (cyan contour). In both figures it is clear that the lift coefficient a is not as well constrained by pitch variation as the drag coefficient b.

---

[12] The grid solution frequency varies as the deployment unfolds. Initially a search is performed every 4 dives until dive 16, when it increases to every 8 dives, until dive 40, when it increases to 16 dives. More frequent early solutions quickly characterize the vehicle, permitting the pilot to adjust onboard flight parameters and minimizing later reprocessing.

[13] The number of data points per-dive can be substantial if the vehicle uses a SciCon science control and data collection module to collect data at high frequency or, in the case of a Deepglider, the dive is especially long and deep.

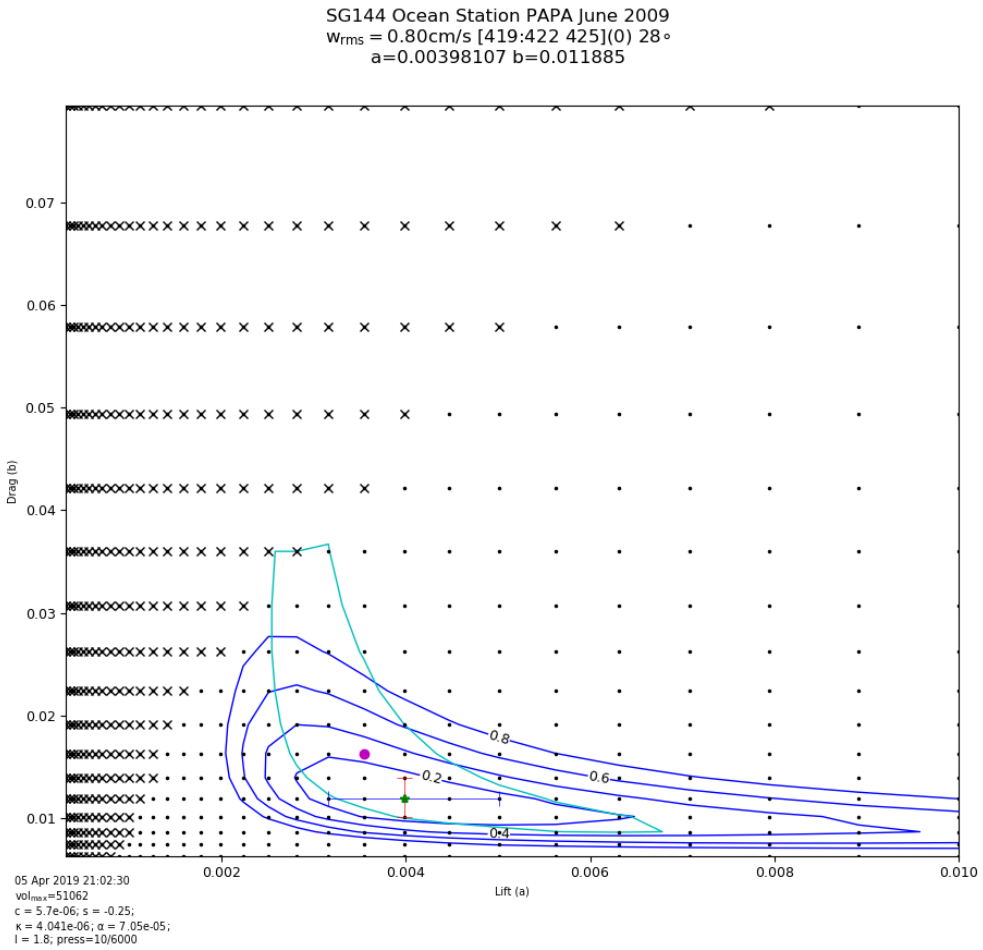

Figure 6: Lift/drag grid solution of w_{rms} combining dives 419 thru 422 and dive 425 from SG144 off Ocean Station PAPA in June 2009, providing 28° pitch separation. Minimum w_{rms} is 0.80 cm/s at the green dot. Contour lines show increases of w_{rms} of 0.2 cm/s. Vertical and horizontal bars indicate w_{rms} solutions within 0.2 cm/s centered on the minimum point. Excessive stalls are indicated by a black x. The red dot and cyan contour indicates the minimum w_{rms} solution and 0.2 cm/s contour from the immediately prior grid solution shown in Figure 5, allowing easy comparison of solutions.

Trusted solutions for *b* are possible only when recent dives have sufficient pitch spread, which depends on how the pilot instructs the vehicle to dive. For long transects between distant targets the desired pitch on sequential dives is often constant; steep dives occur when attempting to achieve nearby targets. If the FMS is unable to determine a trusted solution for two grid-solution periods it alerts the pilot, recommending a steep dive. This can often be accomplished automatically on long transects by placing intermediate targets along the transects. Bowtie sampling patterns, in contrast, provide steep dives frequently enough to avoid the alert request.

FMS provides a summary figure for the *a/b* grid solution trends (see Figure 7). Flying a bow-tie pattern, with regular steep dives at the frequent target points, ensured that nearly all the early drag solutions for SG144 off PAPA are well-constrained. After dive 660, however, during her long transit to the Washington coast, the dive pitches were fairly uniform leading to less well-constrained drag solutions. Nevertheless, the FMS-recovered lift and drag values are fairly consistent over the entire deployment with mean lift coefficient of 0.037 and mean drag coefficient of 0.0133; these compare well with the pilot-supplied values of 0.037 and 0.0120 at the start of the deployment.

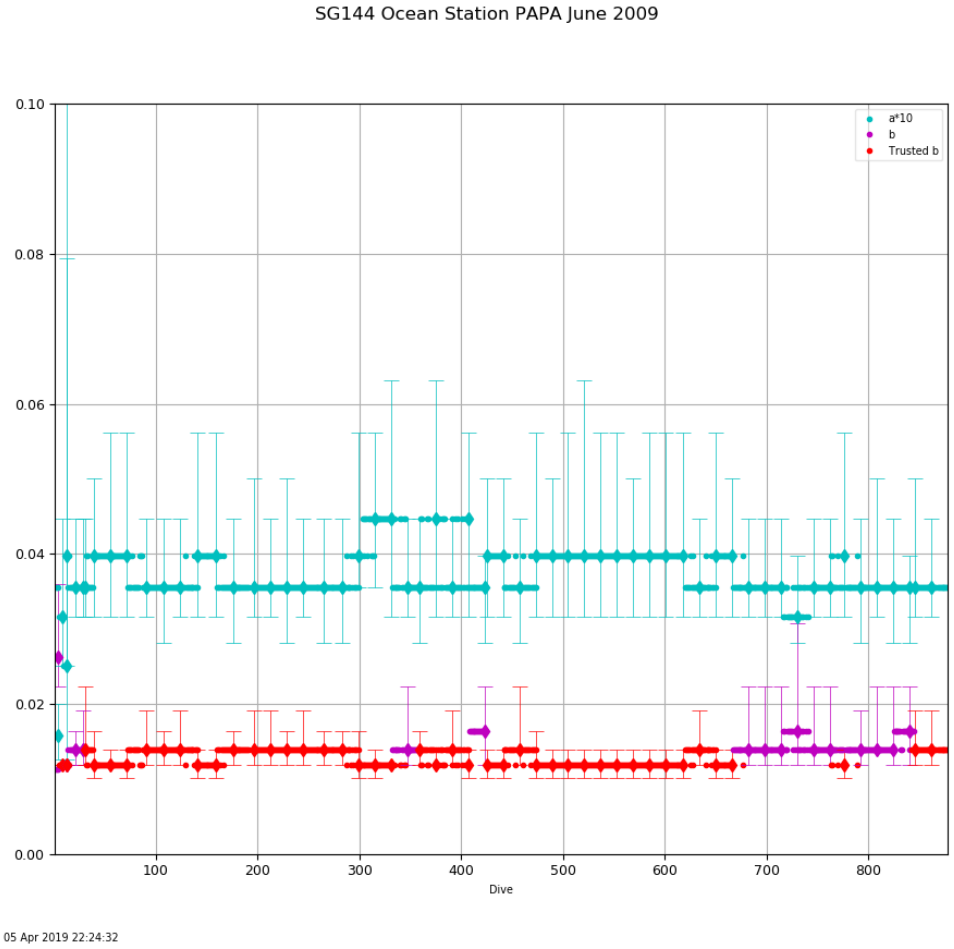

Figure 7: Per-dive lift and drag coefficient solutions for SG144 off Ocean Station PAPA in June 2009. Grid solutions are shown with diamond markers accompanied by vertical bars indicating w_{rms} solutions within 0.2 cm/s; smaller bars (and trusted markers) indicate better constrained solutions for drag.

Results

We ran FMS on 151 deployments over multiple Seaglider types, comprising nearly 54,000 dives. We used the results from these analyses to compute improved default

parameters for each vehicle type, which are reported in an appendix below. Here we discuss FMS results from several different deployments of interest.

Determining the accuracy of FMS predictions using an onboard velocimeter

Two Deepgliders, DG043 and DG044, were fitted with a nose-mounted Rockland Scientific MicroRider turbulence sensor and an RSI AEM1-G-ZR electro-magnetic current velocimeter mounted along the vehicle's primary axis as shown in Figure 8. The velocimeter has a range of 0 to 5 m/s, a resolution of 0.2 cm/s and an accuracy of 0.5 cm/s in our typical speed range of 50 cm/s. Several test dives for both vehicles were flown in Puget Sound to depths of 200 meters.

Figure 8: Velocimeter mounted on the nose of DG043, above the RSI MicroRider turbulence sensor.

The incident water velocity U_{velo} along the glide angle θ is related to the velocimeter-measured velocity $\widehat{U}_{velo}$ as:

$$U_{velo} = \hat{U}_{velo}/\cos\alpha \quad (11)$$

As described, FMS normally minimizes the RMS difference between (pressure sensor) measured and predicted vertical velocity *w*. When velocimeter data is present, FMS minimizes the RMS differences between the model's predicted horizontal and vertical velocities against the normal pressure sensor-based *w* as well as the vertical and horizontal components of the velocimeter speed U_{velo} along the model's predicted glide angle θ.

$$w_{velo} = U_{velo}\sin\theta \quad (12)$$
$$u_{velo} = U_{velo}\cos\theta \quad (13)$$

The lift/drag grid results from the usual *w*-only minimization and then combined with velocimeter data for DG043 are shown in Figure 9.

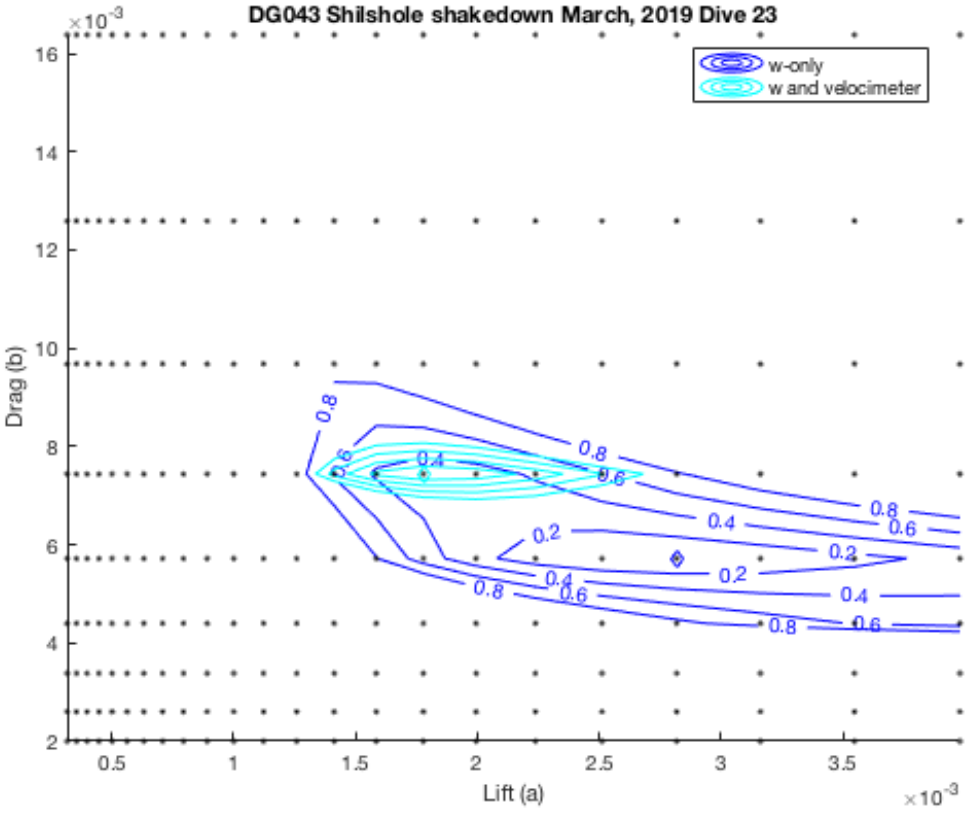

Figure 9: Detail of *w*-only and *w* and velocimeter lift/drag grid solutions combining dives 16 through 23 during the Puget Sound deployment of DG043 in March, 2019. Dives provided 23° pitch separation. Note change of axes scale.

The high pitch separation of the dive set ensures the *w*-only solution is well-constrained in drag but not in lift, as we have observed. However, the inclusion of the velocimeter data, which constrains both the horizontal and vertical velocity

estimates, significantly constrains both the expected drag and lift values. Further, while the *w*-only 0.4 cm/s solution envelope does contain the velocimeter values, the minimum points are different: The *w*-only solution tends to under-estimate the vehicle drag by about 30 percent (0.00572 vs. 0.00744) and to over-estimate lift by about 60 percent (0.001585 vs. 0.000891) compared to the velocimeter-based solution.

Running the same procedure for DG044 deployed in Puget Sound in August, 2018 shows a similar pattern (Figure 10). Again, the *w*-only procedure apparently underestimates the vehicle drag by about 30 percent (0.00744 vs. 0.00968) compared to the velocimeter solution. Note that the absolute drag coefficients are overall higher for DG044 than for DG043 (using both solution types); this is due to DG044 carrying a WETLabs ECO sensor, which protrudes from the aft fairing.

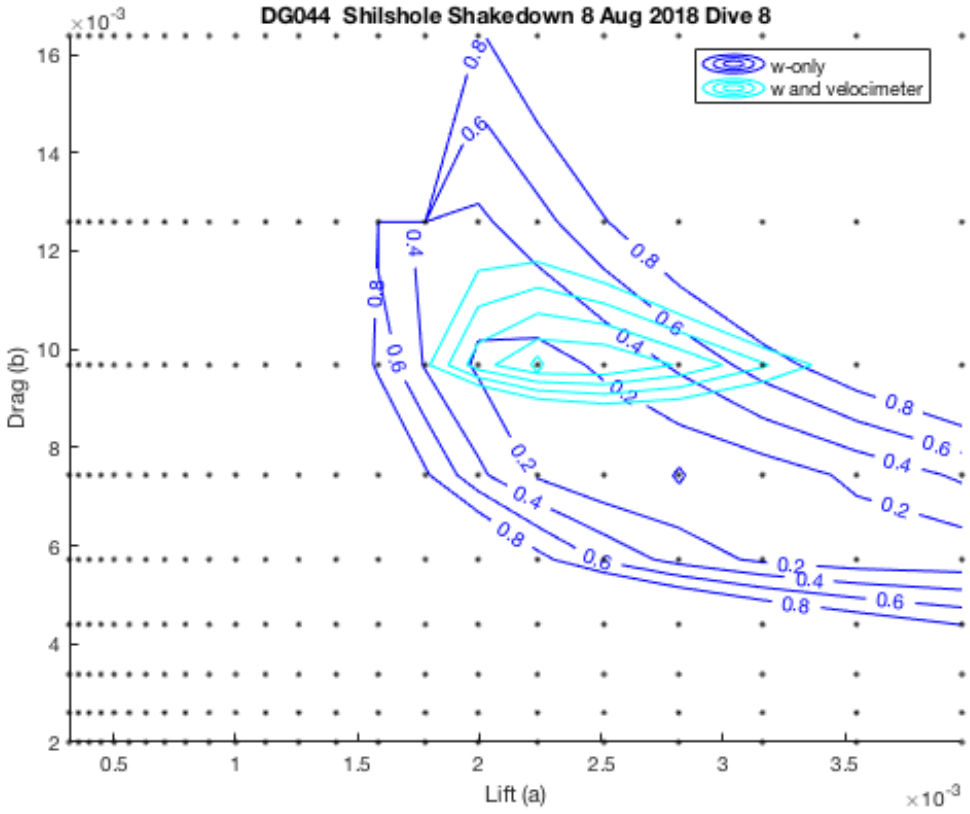

Figure 10: Detail of *w*-only and *w* and velocimeter lift/drag grid solutions combining dives 3, 5, 7, and 8 during the Puget Sound deployment of DG044 in August, 2018. Dives provided 16° pitch separation. Note change of axes scale.

Analysis of a subsequent deployment of DG044 off Iceland in June, 2019 diving to 900 meters shows that, using *w*-only, FMS similarly under-estimates drag and over-estimates lift compared to incorporating velocimeter data; see Figure 11. While the drag coefficient values are quite similar to those found during the Puget Sound

deployment the lift coefficient is substantially smaller; this is due to the consistently steep (45°), fast (50 cm/s) dives performed off Iceland but not in Puget Sound.

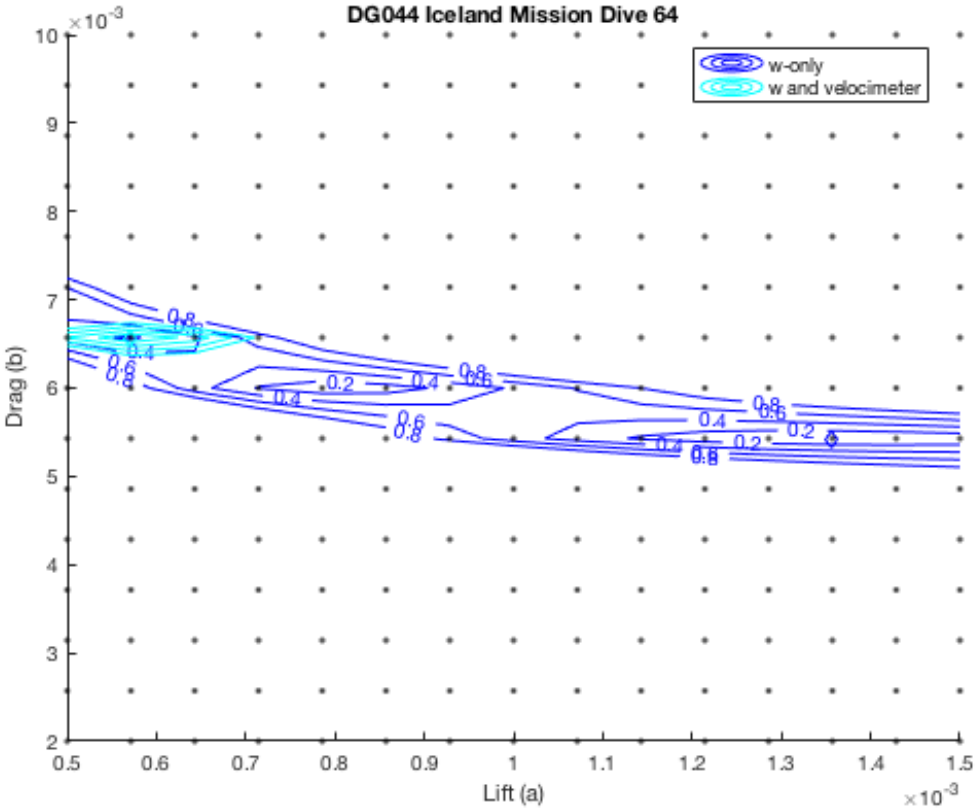

Figure 11: Detail of *w*-only and *w* and velocimeter lift/drag grid solutions combining dives 48 thru 56 and 58 thru 64 during the Iceland deployment of DG044 in June, 2019. Although these steep dives provided only a 5° pitch separation, drag is nevertheless well constrained for both solutions. Note that for this analysis a finer grid was used to estimate the lift/drag coefficients. Note also that, in this case, there are two potential *w*-only w_{rms} minima, unlike the single *w* and velocimeter solution.

The *w*-only grid solution is overfitting the *w* data as shown in the following extended vertical-velocity plots for dive 22 from DG043 in March, 2019 using both the *w*-only and the velocimeter-constrained solutions. The vertical velocity residuals are smaller based on the *w*-only solution (Figure 12) but the horizontal speed predictions are substantially worse than the velocimeter-estimated values. Figure 13 shows that while the vertical velocity residuals are slightly worse, the horizontal velocity estimates are much better when employing velocimeter-informed lift/drag coefficients.

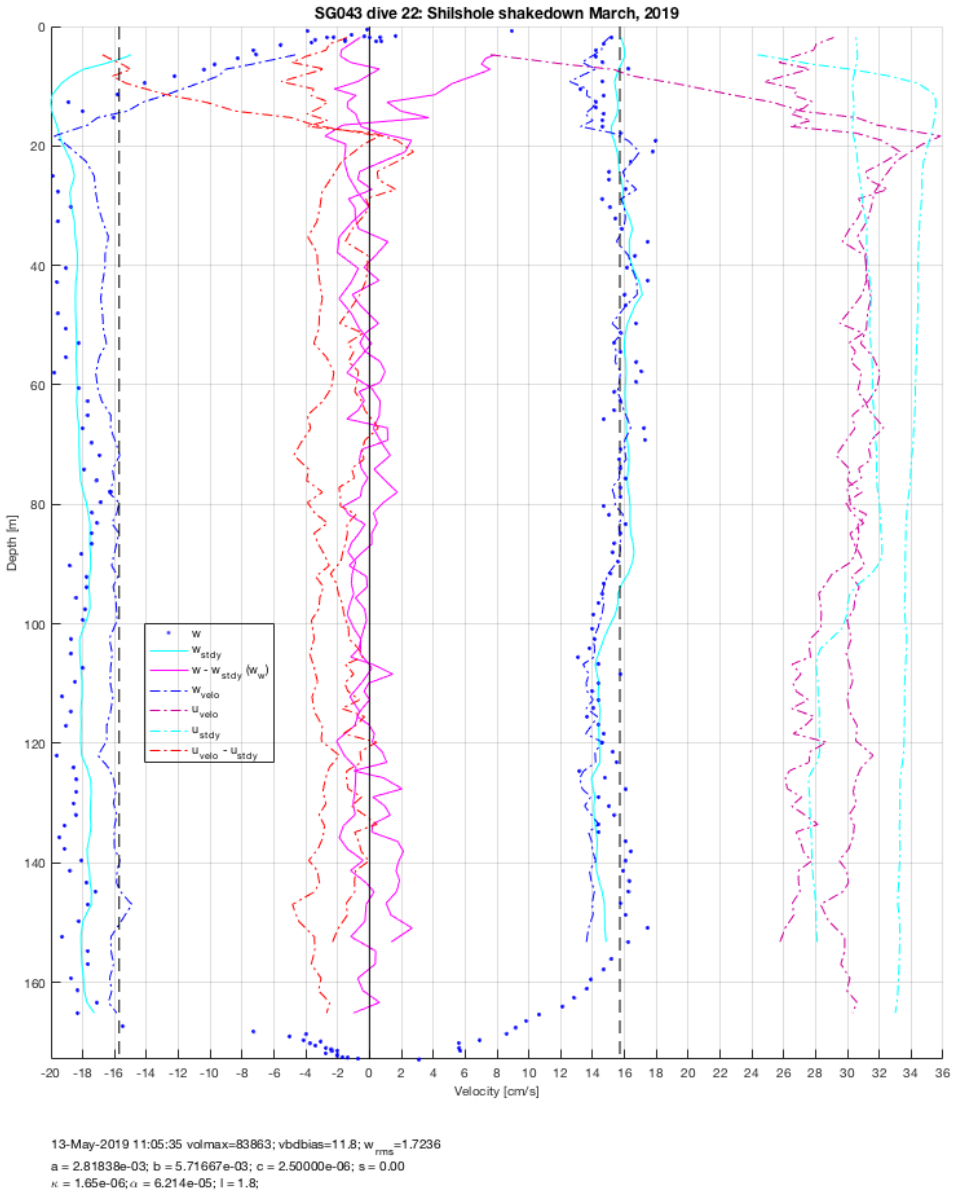

Figure 12: Vertical and horizontal velocity plot for dive 22 from DG043 in March, 2019 using the w-only lift/drag solution. Vertical dashed lines indicate the desired w. Note the poor horizontal velocity u_{stdy} prediction, especially on the dive, which overestimates u_{velo} by ~4 cm/s.

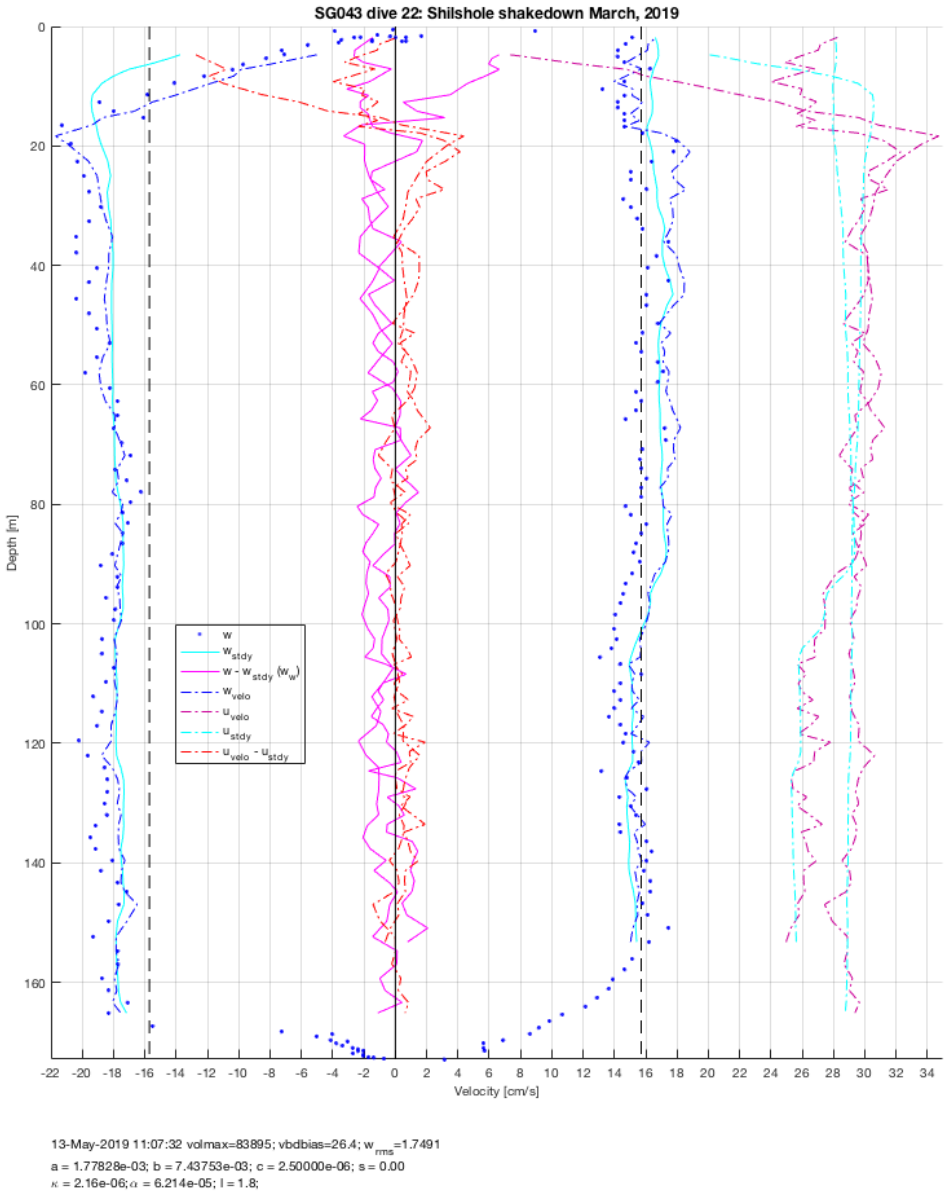

Figure 13: Vertical and horizontal velocity plot for dive 22 from DG043 in March, 2019 using the combined *w* and velocimeter lift/drag solution. Vertical dashed lines indicate the desired *w*. Note the improved horizontal velocity fit.

Figure 14 below shows the impact of scaled lift and drag coefficients on the depth-averaged current (DAC) estimates for dive 13 from DG043 in March, 2019. The velocimeter-based drag estimate is 30% larger than the minimum shown in the figure. This increased drag value corresponds to the grid entries just above the minimum point. The cyan magnitude contours (and directional arrows) indicate velocimeter-informed drag value would yield a DAC magnitude change of about 2

cm/s to the southwest[14]; this is a 10% difference on the 21 cm/s depth-average current estimated for this dive[15]. Note that differences in scaled lift have little impact on the estimated depth-averaged current.

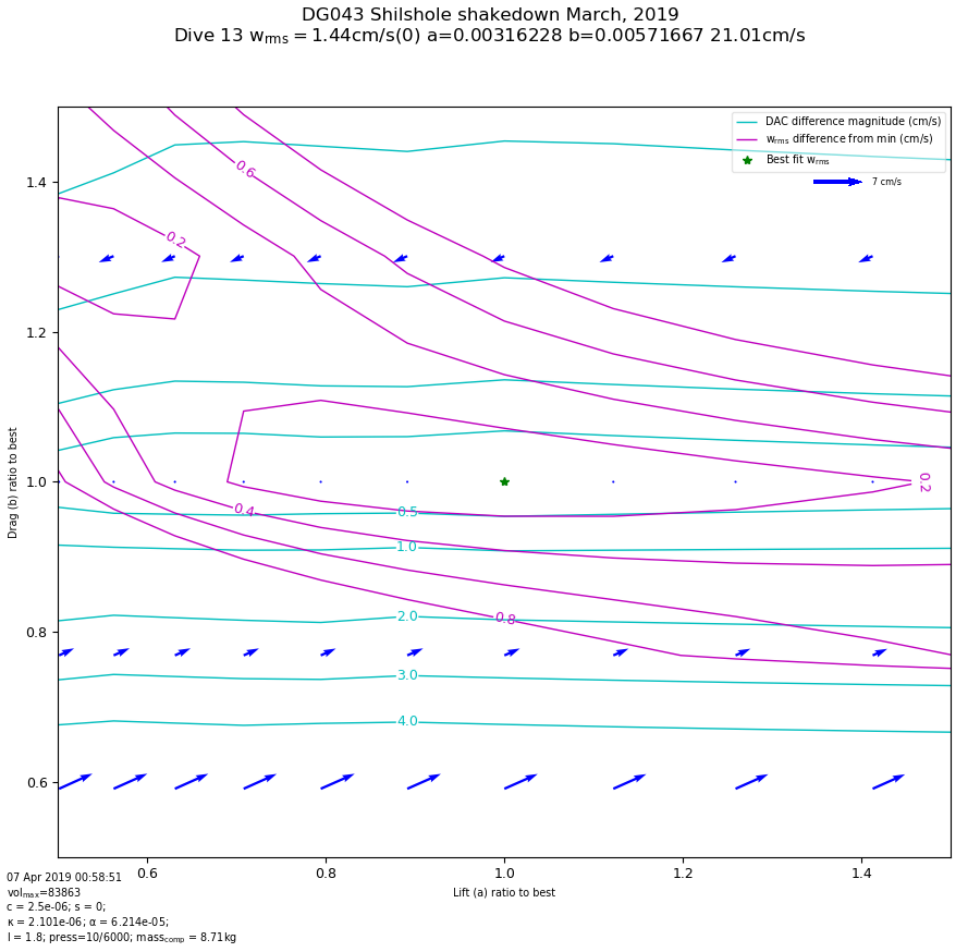

Figure 14: Impact of scaled lift and drag coefficients on depth-averaged current (DAC) estimations for dive 13 of DG043 in March, 2019. Lift and drag coefficients are shown scaled from the minimum of a and b established using w-only comparisons. Magenta contours show the w_{rms} difference contours shown in Figure 9 appropriately scaled. Blue arrows show the magnitude and direction of DAC *differences* from the 21 cm/s DAC estimated at the minimum lift/drag point (green star). Cyan contours show the DAC magnitude differences.

[14] The overall direction of the depth-average current for a dive remains the same as lift and drag are scaled, as can be seen in Figure 14. Changes in lift and drag change the magnitude of the horizontal velocity component for each data point but the direction is given by that point's compass heading, which is unchanged. Overall the combined displacement vectors therefore simply change the DAC magnitude.

[15] The overall accuracy of DAC estimates also depends upon heading accuracy from the calibrated compass and the accuracy of the GPS positional measurements at dive and surfacing. See Bennett, et al. 2014 for a discussion about historical GPS accuracy on Seagliders.

Systematic over/under-estimates of the lift and drag coefficients, respectively, by the *w*-only processing impacts the DAC estimated on each dive and thus could impact large-scale analysis of ocean transports and flows. We assessed the impact on DAC analysis of a heuristic adjustment of the recovered *w*-only lift and drag parameters on the nine-month deployment of SG128 off the Washington (USA) coast in 2015 as part of the Northwest Association of Networked Ocean Observing Systems (http://www.nanoos.org/). In this modified reprocessing, the drag coefficient was increased by 20% (1.2 x b) and the lift coefficient decreased by 40% (0.6 x a).

Figure 15 and Figure 16 show plots of components of DAC estimated under the "improved" heuristic adjustment versus those under the "original" *w*-only solution for each dive of the deployment. For each dive, the DAC is broken into components parallel to (Figure 15) and orthogonal to (Figure 16) the net through-water displacement. The DAC component orthogonal to the through-water displacement is almost completely unaffected by the heuristic adjustment (Figure 16). The component parallel to the through-water displacement in the "improved" processing is highly correlated with the "original" but has a mean positive offset of 1.9 cm/s (Figure 15). In other words, after adjustment, the average vehicle speed is 1.9 cm/ s slower—hence DAC is greater—along the net through-water movement direction, consistent with the increased drag and decreased lift parameters.

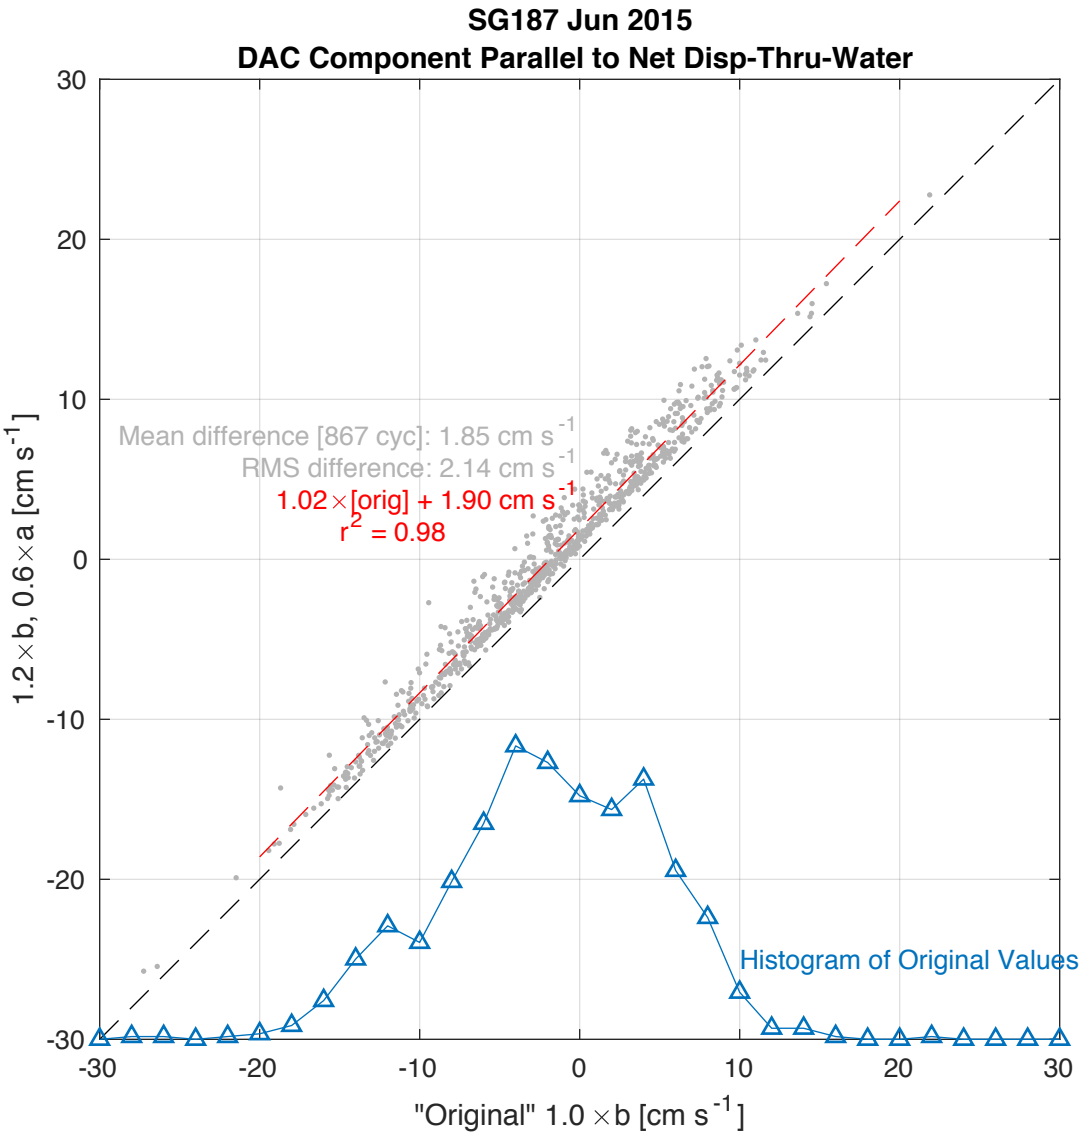

Figure 15: Comparison of DAC components for SG128's June 205 deployment parallel with net thru-water displacement between *w*-only ("original") and heuristically-adjusted *w*-only ("improved") solutions.

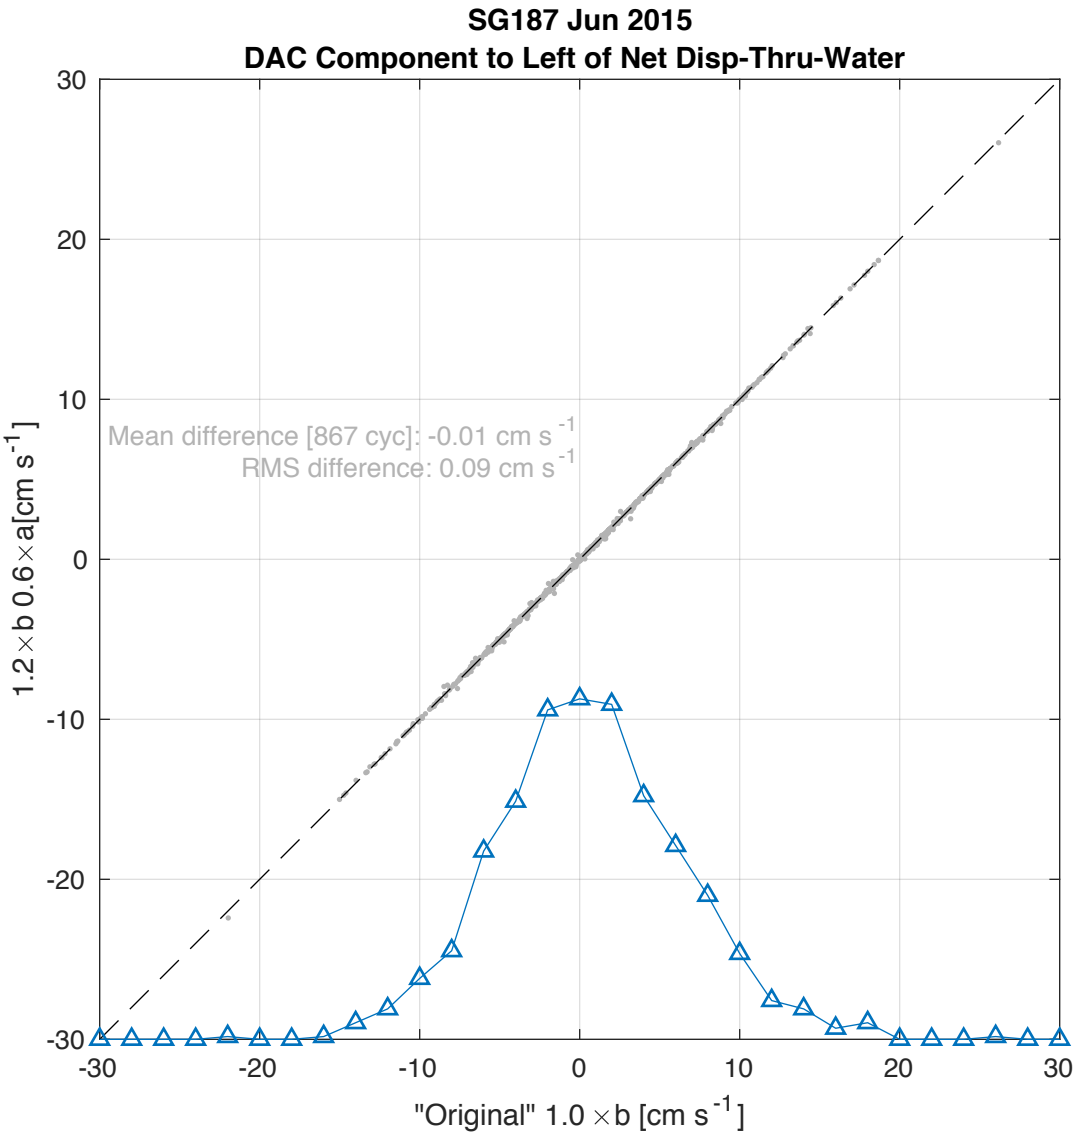

Figure 16: Comparison of DAC components for SG128's June 205 deployment orthogonal to net thru-water displacement between *w*-only ("original") and heuristically-adjusted *w*-only ("improved") solutions.

One of the most common uses of DAC is as a reference velocity for estimates of depth-dependent geostrophic flows. The effect that a systematic along-track offset in DAC might have on reference velocity depends on the sampling configuration. In common configurations such as coastal line transects or open-ocean repeat survey patterns, the effect is likely to be small with respect to typical oceanic mesoscale flows. In the survey pattern case, estimating the reference velocity over multiple repeats of the pattern (e.g., Pelland et al., 2016 at Ocean Station PAPA) results in error cancellation as the vehicle travels across the symmetric pattern. In the line transect case, the relevant component of the reference velocity is generally orthogonal to the vehicle track and results above show that this component of DAC

is unaffected. In the SG187 2015 example, the DAC data from cross-shore line transects performed by this vehicle were mapped to a regular grid to estimate reference velocity following methods described in Pelland et al., 2013. This mapping was computed for both the adjusted and unadjusted processing versions. The mapped reference velocity, subsampled in 5 km increments (n=738), differed on average by -0.13 cm/s (adjusted minus unadjusted, positive alongshore) across the deployment. The maximum difference was 2.26 cm/s, and 89.6% of differences were <1 cm/s in magnitude. These results support the idea that in the line transect case, the utility of DAC as a reference velocity is not likely to be significantly degraded by lift and drag adjustments similar in character to those of the velocimeter-constrained Deepglider solutions. However, additional data from velocimeters (or vehicle-mounted ADCP sensors) on different vehicle types are needed before any heuristic scaling adjustment to the *w*-only solutions can be recommended.

The 2004 Labrador Sea Deployments

Figure 17, Figure 18, and Figure 19 summarize the results of FMS on the 2004 Labrador Sea deployment analyzed by FW2011.

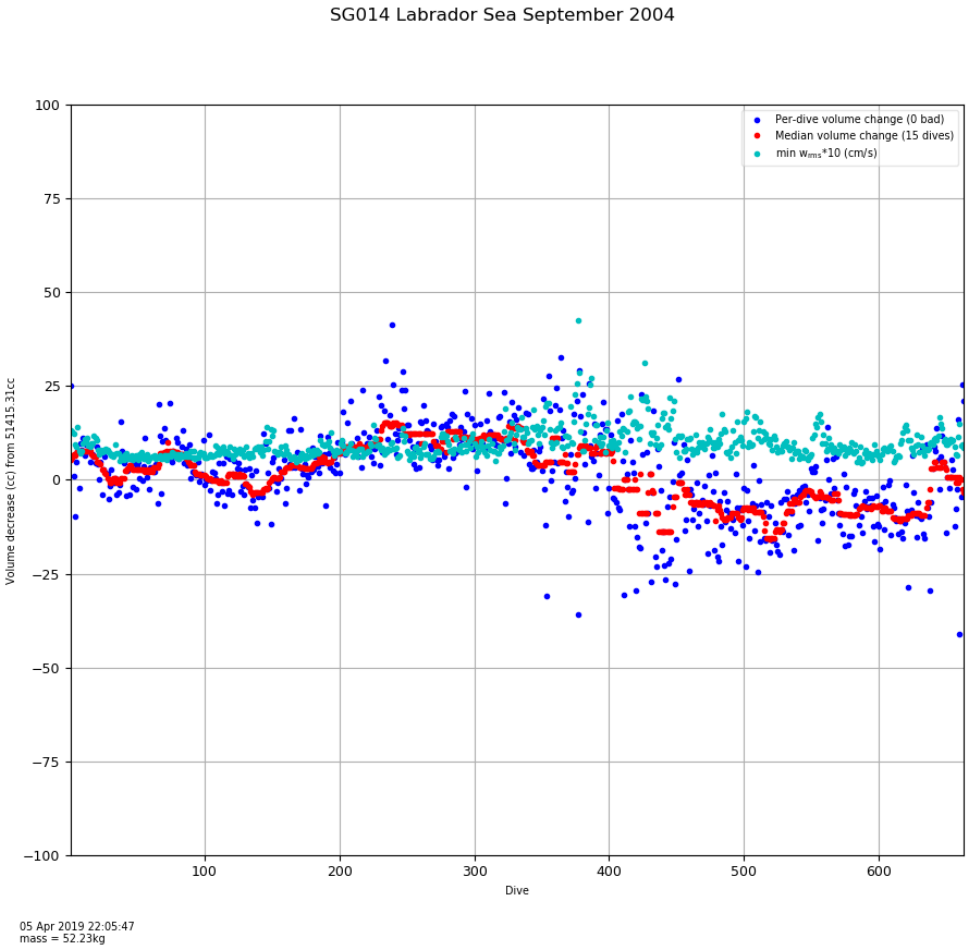

Figure 17: Per-dive solutions of δV_0 for SG014 in the Labrador Sea in 2004. V_0 was determined to be 51415 cc. Blue points are per-dive estimates. Red points are a 15-dive median filtered trend. Cyan points show scaled per-dive w_{rms} values.

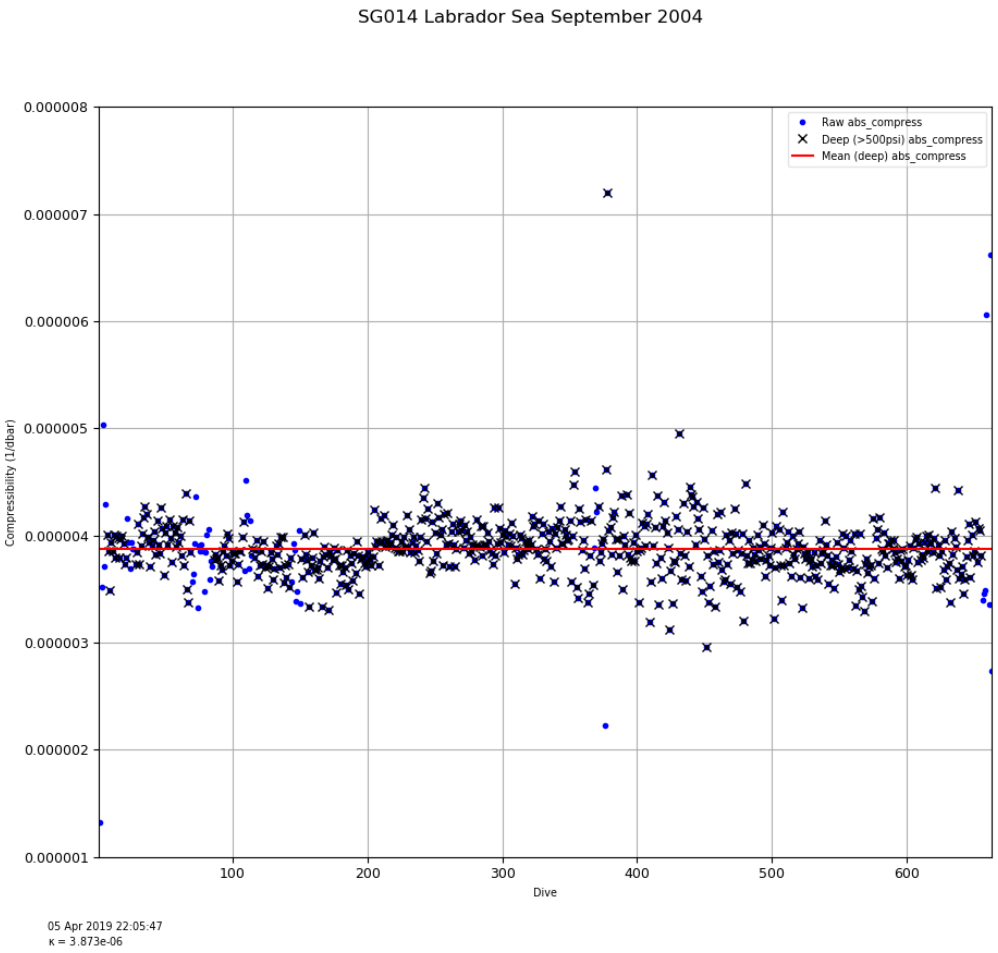

Figure 18: Per-dive solutions of κ for SG014 in the Labrador Sea during 2004. Early dives did not reach 1000m (see Figure 20 below). Mean κ was determined to be $3.873e^{-6}$.

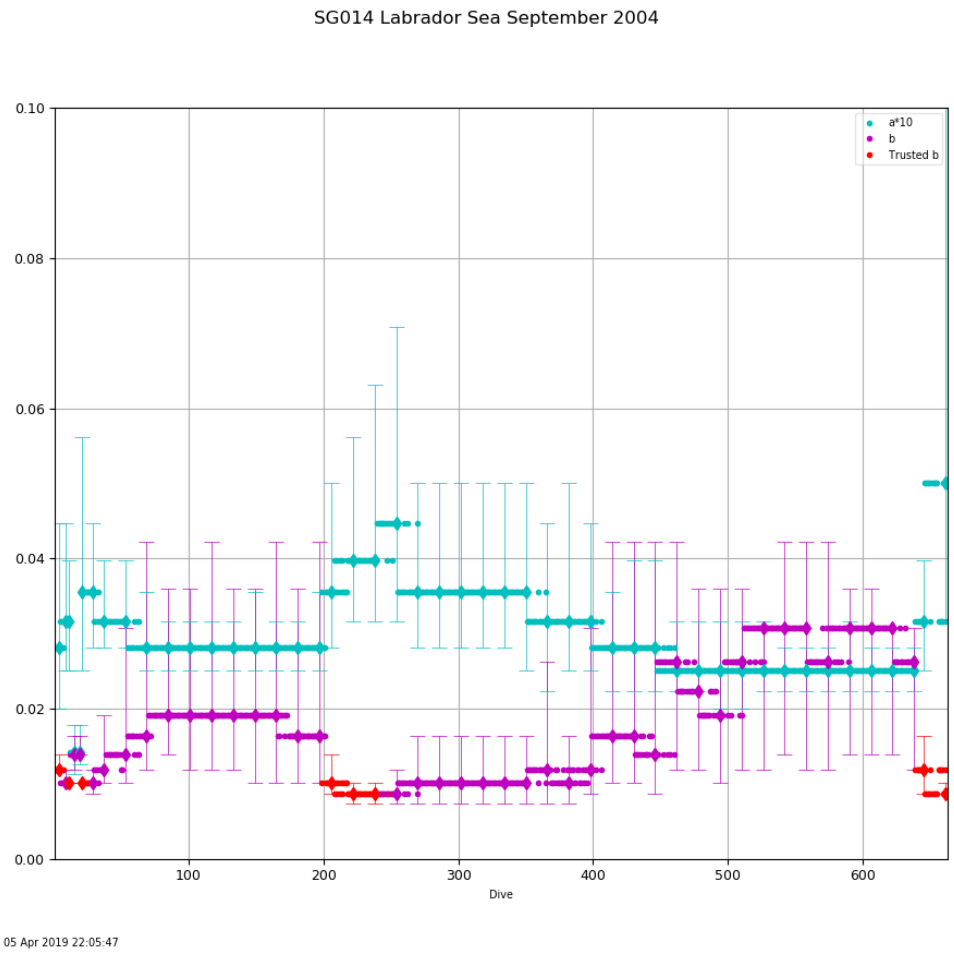

Figure 19: Per-dive lift and drag solutions for SG014 in the Labrador Sea during 2004. Grid solutions are shown with diamond markers; vertical bars indicate w_{rms} solutions within 0.2 cm/s; smaller bars (and trusted markers) indicate better constrained solutions for drag.

Overall, FMS recovered a V_0 value of 51415 cc, close to the value found by FW2011: 51400 - (-12.4) = 51412.4 cc (their Equation 10). However, Figure 17 shows the per-dive δV_0 is not constant over the deployment with an apparent volume increase after dive 400. At roughly the same dive Figure 19 shows a gradual increase in drag and decrease in lift, characteristic of potential biofouling.

FW2011 employed the first 50 dives to determine a single set of lift and drag coefficients to characterize the entire deployment. FMS performed several lift/drag grid searches before dive 50; those mean values are, within the resolution of our fixed grid scheme, very close to the values FW2011 found: a=0.0028 vs. their 0.0040 and b=0.0118 vs. their 0.0088. Figure 18 shows an overall *mean* κ from dives deeper than 500psi at 3.783e-6, smaller than the 4.1e-6 value FW2011 reported; if the mean is restricted to the first 50 dives, the FMS value is slightly higher at 3.882e-6.

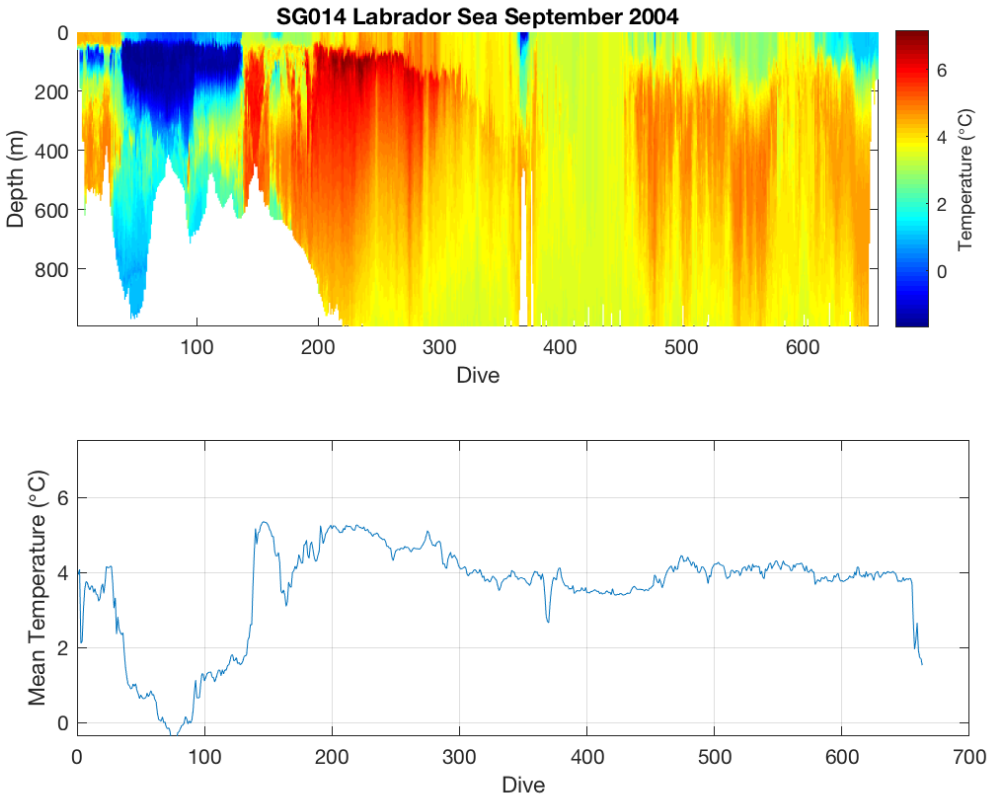

Figure 20: Seawater temperature recorded by SG014 in the Labrador Sea in 2004. Mean temperature is shown in the bottom panel.

Figure 20 reports the temperature against depth from SG014 during this deployment. We note that the initial excursion of lift and drag coefficients found by FMS after dive 50 roughly corresponds to entering fairly cold waters during relatively shallow dives. However, there is no apparent change in volume or compressibility likely ruling out thermal shrinkage of oil in the external VBD bladder.

FW2011 did not provide the results of applying their procedure to SG015's 2004 deployment in the Labrador Sea. For the record we present the FMS solutions below; see Figure 21 thru Figure 23. This vehicle also experienced an increase in volume after starting to dive deep; however, the suggestion of potential biofouling is delayed until after dive 500.

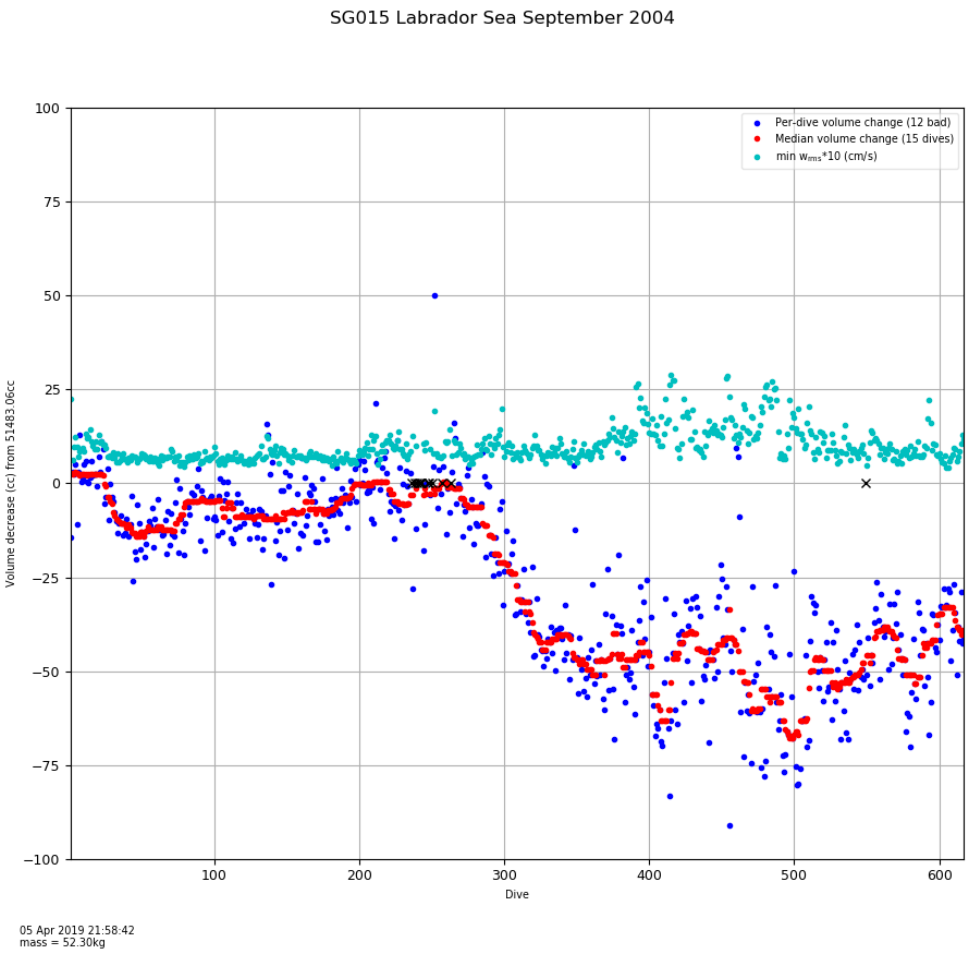

Figure 21: Per-dive solutions of δV_0 for SG015 in the Labrador Sea in 2004. V_0 was determined to be 51483 cc. Blue points are per-dive estimates. Red points are a 15-dive median filtered trend. Cyan points show scaled per-dive w_{rms} values.

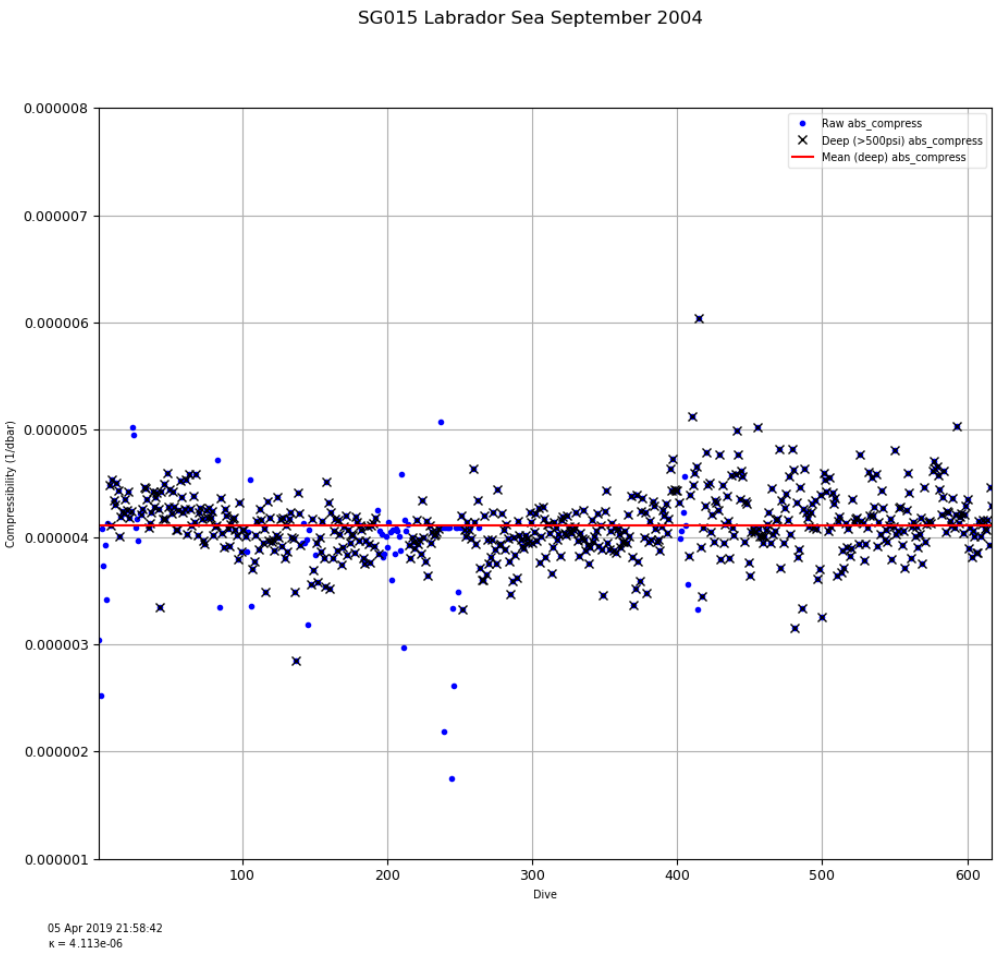

Figure 22: Per-dive solutions of κ for SG015 in the Labrador Sea during 2004. Early dives did not reach 1000m (see Figure 24). Mean κ was determined to be $4.113e^{-6}$, higher than SG014.

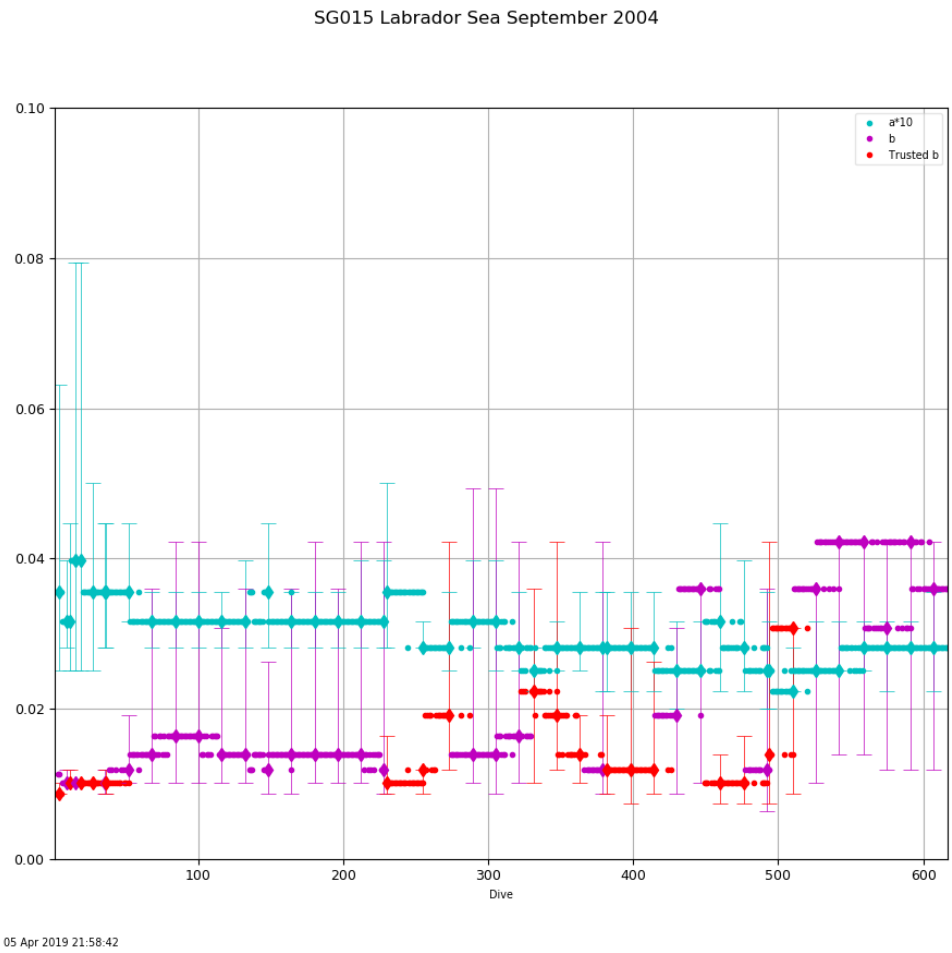

Figure 23: Per-dive lift and drag solutions for SG015 in the Labrador Sea during 2004. Grid solutions are shown with diamond markers; vertical bars indicate w_{rms} solutions within 0.2 cm/s; smaller bars (and trusted markers) indicate better constrained solutions for drag. Note the substantial increase in drag after dive 500 without a decrease in lift, although there is an overall decreasing trend on lift over the deployment. In addition, there is a temporary rise and fall in (trusted) drag around dive 340; we observed similar transient changes in flight regime in other deployments.

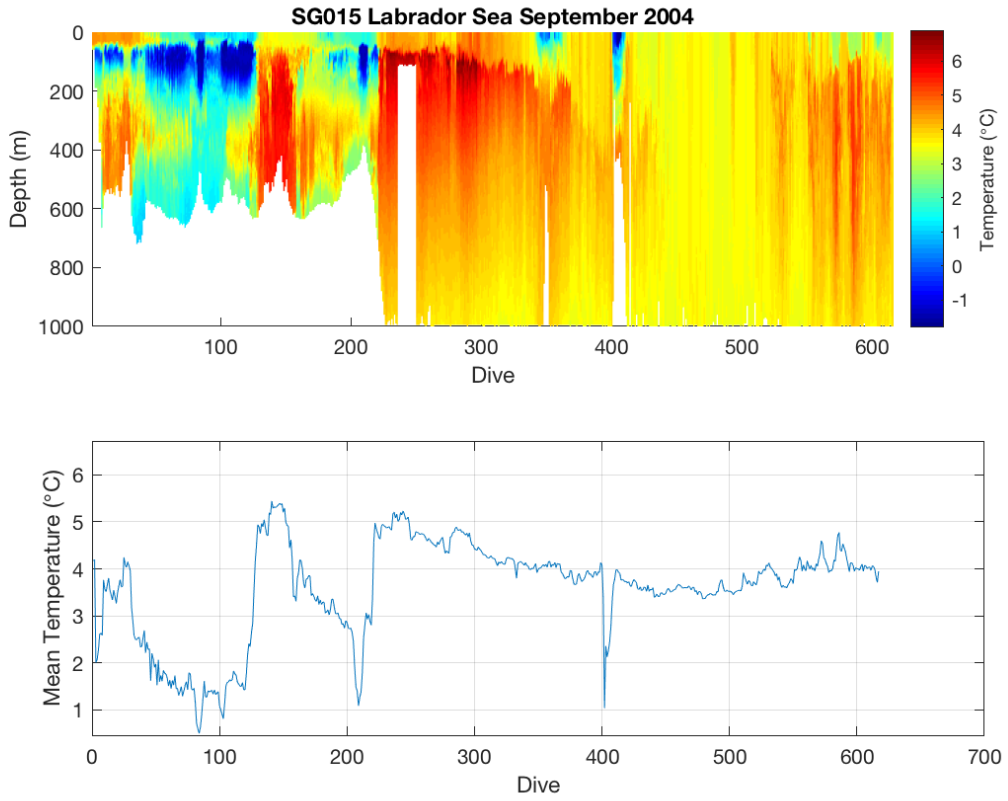

Figure 24: Seawater temperature recorded by SG015 in the Labrador Sea in 2004. Mean temperature is shown in the bottom panel.

Normal and 'abnormal' deployments

We classified deployments as 'normal' if it appeared that a single set of flight parameters could characterize the entire deployment, the original assumption of the basestation processing. In particular, we considered a deployment's flight 'normal' if the difference between the maximum and minimum median-filtered δV_0 was less than 70cc (approximately 0.1% of typical V_0 values) and the standard deviation of b was less than 0.003 (approximately 15% of typical b values). This permitted some gradual decrease in volume due to fairing water saturation and some small variation in b estimation due to overfitting sensor noise.

Under these criteria, we found 104 of 151 (69%) deployments were 'normal' (covering 30,389 dives), while the balance (47 deployments, 31% covering 23,347 dives) demonstrated either substantial volume changes or flight regime changes or both. Under these criteria, the June, 2009 deployment of SG144 off Ocean Station PAPA discussed above is 'normal' but the two 2004 Labrador Sea deployments are not. We discuss some additional abnormal deployments below.

Strong Biofouling in the Southern Ocean

In February and September, 2012, as part of the Southern Ocean Seasonal Cycle Experiment (SOSCEx), two Seagliders SG573 and SG574 were deployed at 42°S, 11°W and made their way east (Swart, 2012). After several weeks of routine flight, the pilots for both deployments reported having trouble maintaining course; the vehicles frequently entered flat-spins. Recovery revealed substantial infestations of goose-neck barnacles on the flight surfaces of the vehicles (see Figure 25).

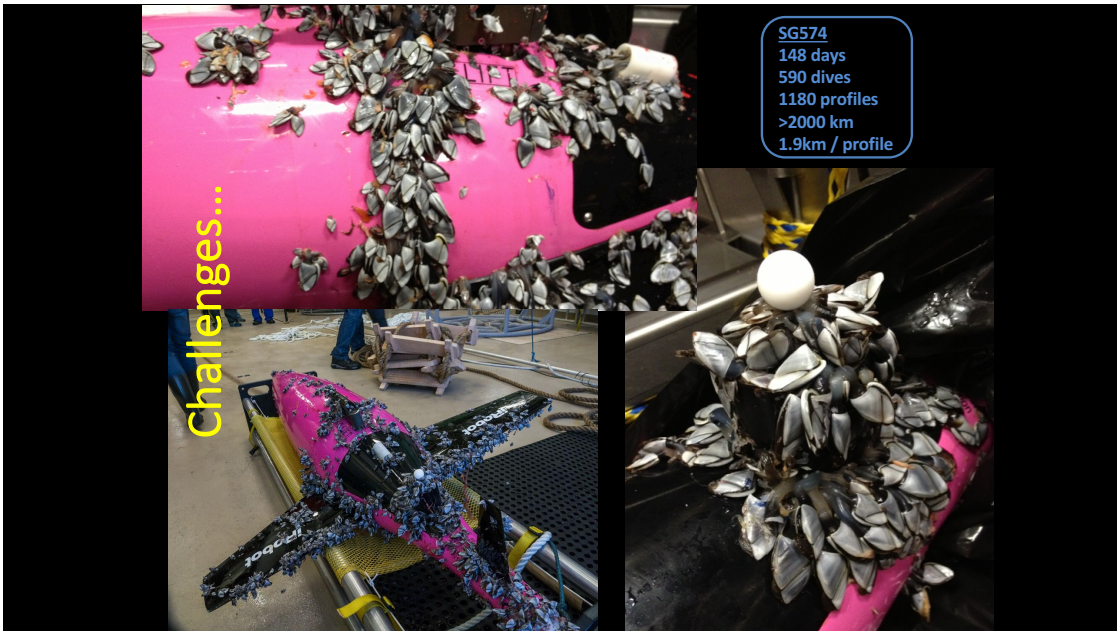

Figure 25: Post-recovery images of SG574 after deployment in the Southern Ocean, 2012 as part of SOSCEx.

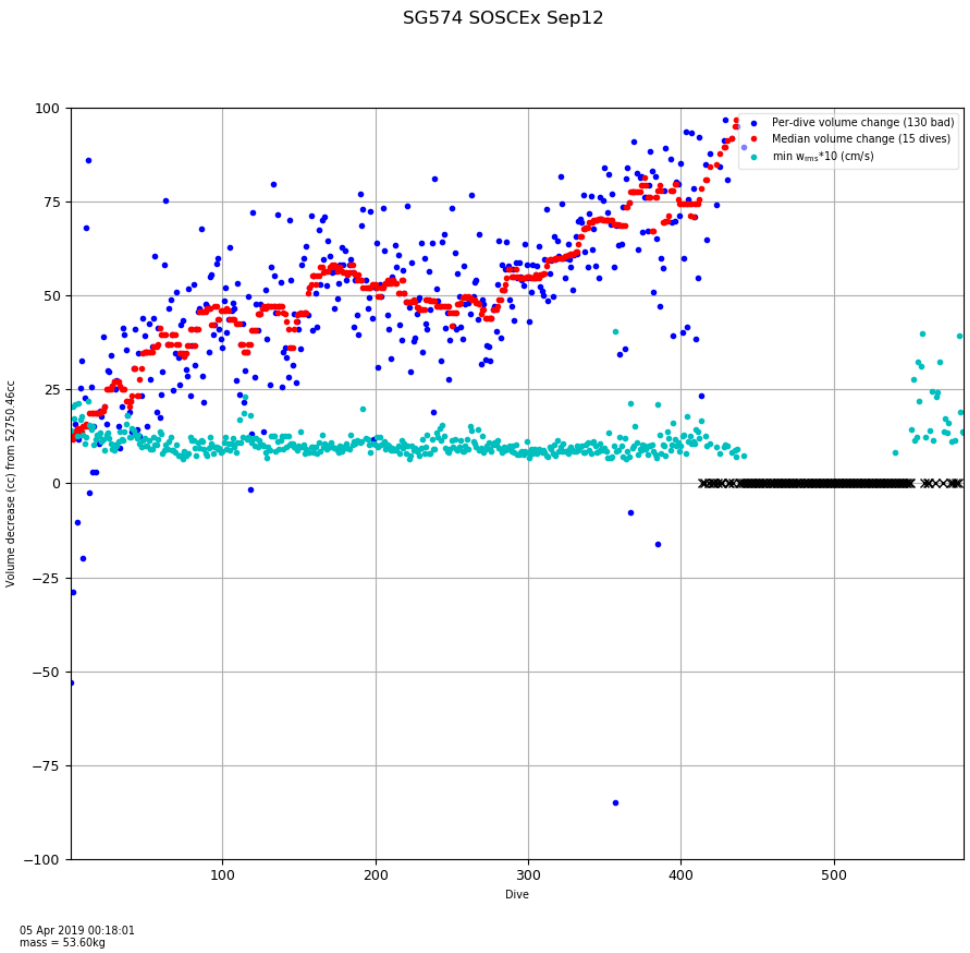

Figure 26: Per-dive solutions of δV_0 for SG574 in the Southern Ocean in 2012. Blue points are per-dive estimates. Red points are a 15-dive median filtered trend. Cyan points show scaled per-dive w_{rms} values.

Figure 26 shows an immediate and substantial decrease in volume that continued throughout the deployment. Indeed, it prohibited FMS from finding consistent solutions within its search bound after dive 410; all these dives reported flat spins. As discussed above, the increase in barnacle mass and volume would change the overall vehicle density. However, since there is no independent way of measuring density change during a deployment, our model reflects any changing density as a changing volume against the assumed fixed mass of the vehicle.

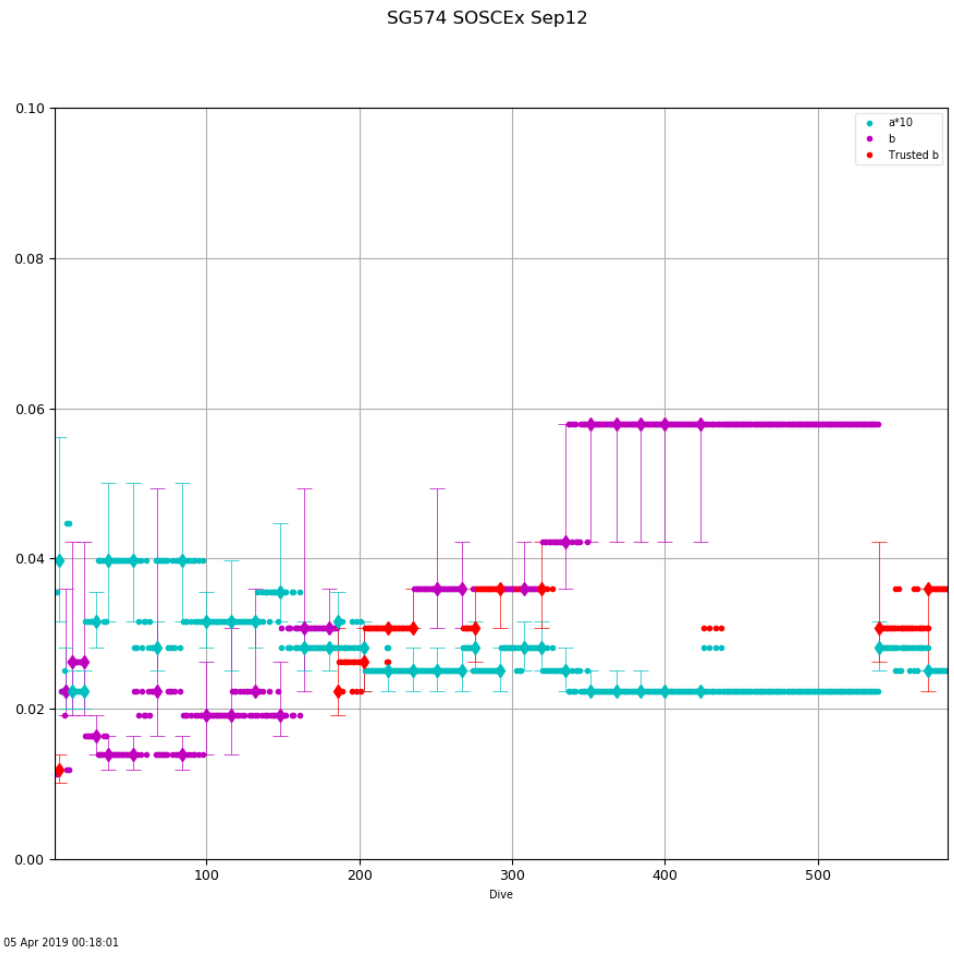

Figure 27: Per-dive lift and drag solutions for SG574 in the Southern Ocean in 2012. Grid solutions are shown with diamond markers; vertical bars indicate w_{rms} solutions within 0.2 cm/s; smaller bars (and trusted markers) indicate better constrained solutions for drag.

Figure 27 shows immediately increasing drag and decreasing lift consistent with a growing number of barnacles. Rudnick, et al., 2013, their Figure 4 and 5, reported similar increases in recovered drag estimates during a deployment of Spray gliders off the Luzon Strait (assuming, it appears, a constant lift coefficient).

Although the volume and lift/drag regimes were rapidly changing Figure 28 shows the absolute compressibility of the vehicle was largely unchanged, as might be expected since it is dominated by the hull and sensors, and the barnacles are roughly as compressible as seawater.

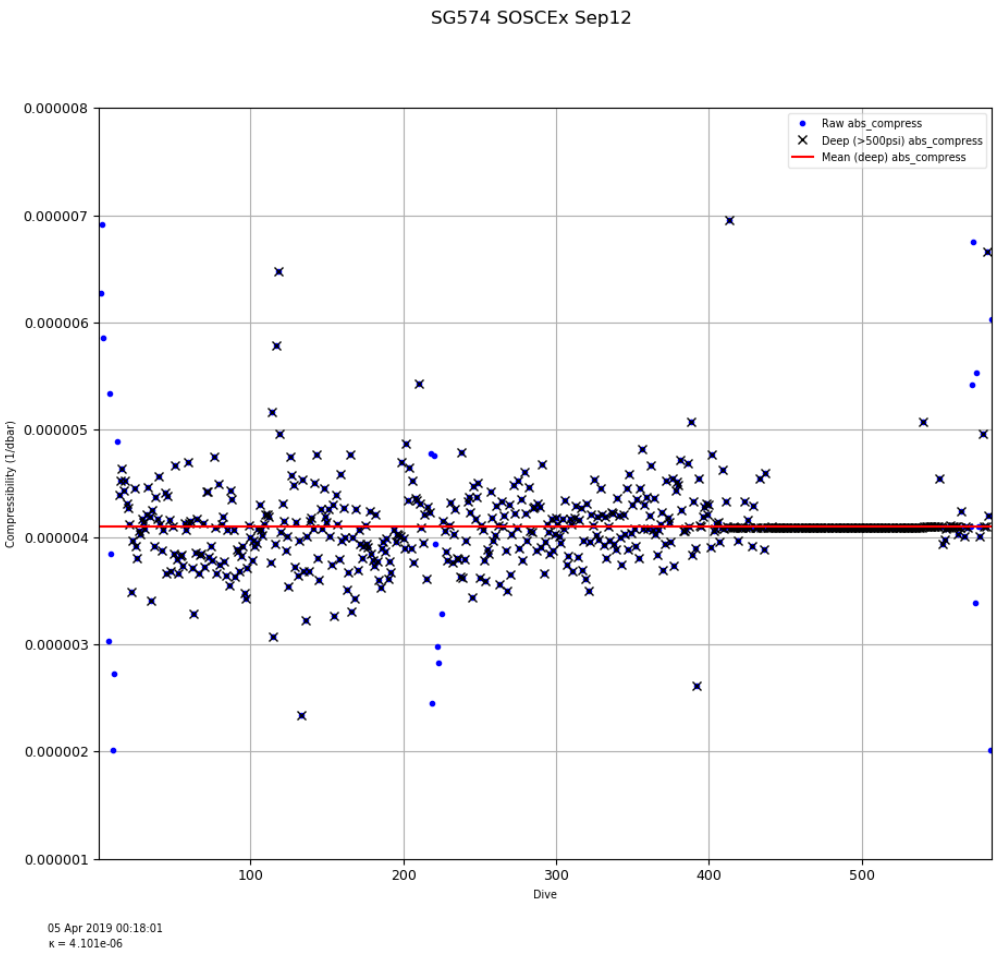

Figure 28: Per-dive solutions of κ for SG574 in the Southern Ocean in 2012.

A similar biofouling incident occurred on the other glider, SG573, also deployed during SOSCEx in the Southern Ocean in 2012.

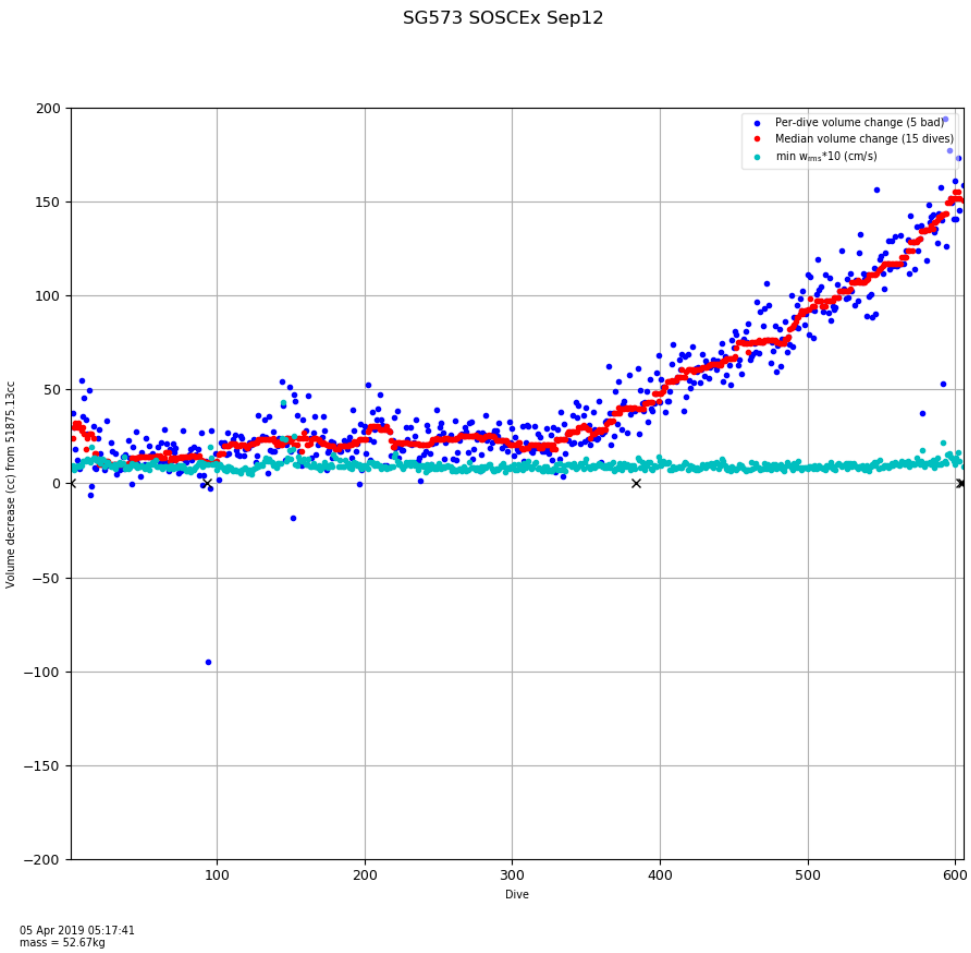

Figure 29: Per-dive solutions of δV_0 for SG573 in the Southern Ocean in 2012. Blue points are per-dive estimates. Red points are a 15-dive median filtered trend. Cyan points show scaled per-dive w_{rms} values.

Figure 29 shows the rapid, apparent decrease of vehicle volume after dive 370. Calculations show the volume of SG573 decreased by 150cc, implying the minimum density of the vehicle increased from 1.0153g/cc to 1.0183g/cc.

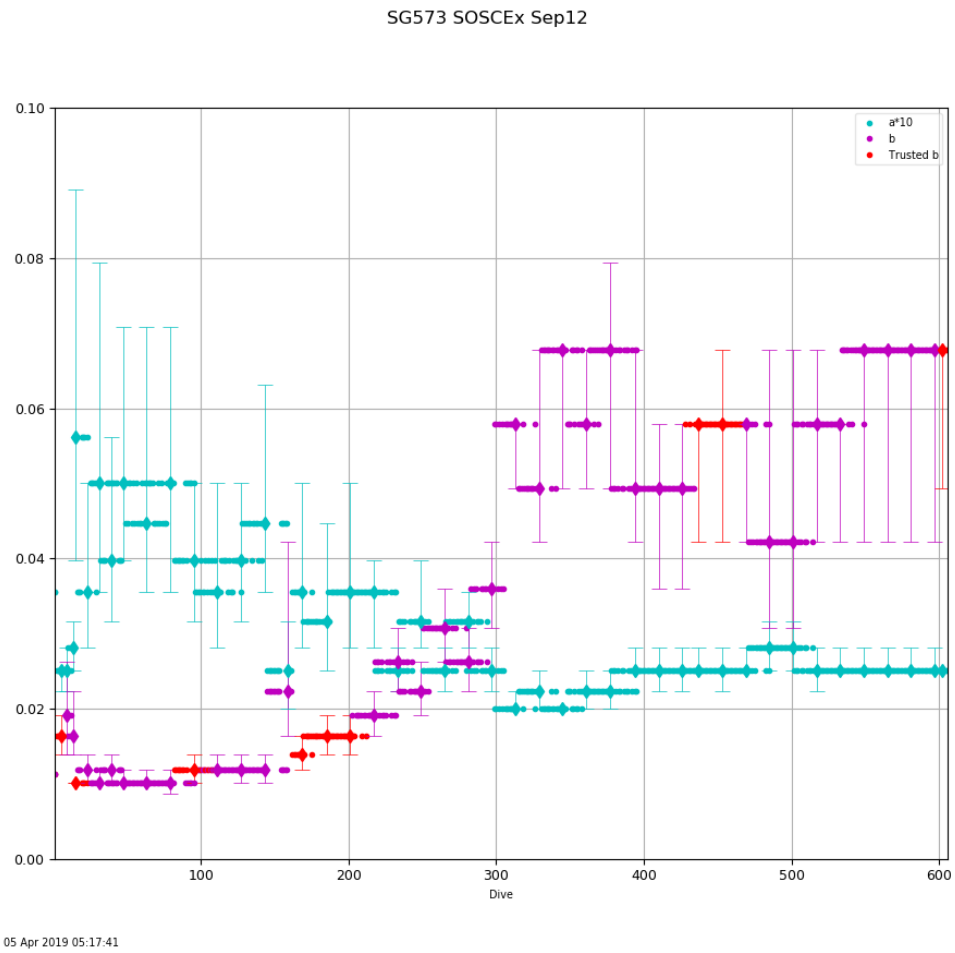

Figure 30: Per-dive lift and drag solutions for SG573 in the Southern Ocean in 2012. Grid solutions are shown with diamond markers; vertical bars indicate w_{rms} solutions within 0.2 cm/s; smaller bars (and trusted markers) indicate better constrained solutions for drag.

Figure 30 shows the increasing drag and decreasing lift parameters beginning after dive 160. By the end of the deployment the drag had increased by a factor of six and the lift had decreased by a factor of four. If we assume the barnacles were the sole flight issue with this glider and that they started growing around dive 160, their impact on volume was not detectable for nearly 200 dives (about 50 days).

Clearly no single flight regime can characterize either of these deployments. However, because FMS determines and applies the proper flight regime for each dive, the predicted vertical velocities for even the most biofouled dives remain accurate and hence derived quantities such as depth-averaged current remain trustworthy. Figure 31 and Figure 32 below demonstrate accurate recovery of vertical velocity estimates for SG573 in two very different flight regimes during her 2012 deployment.

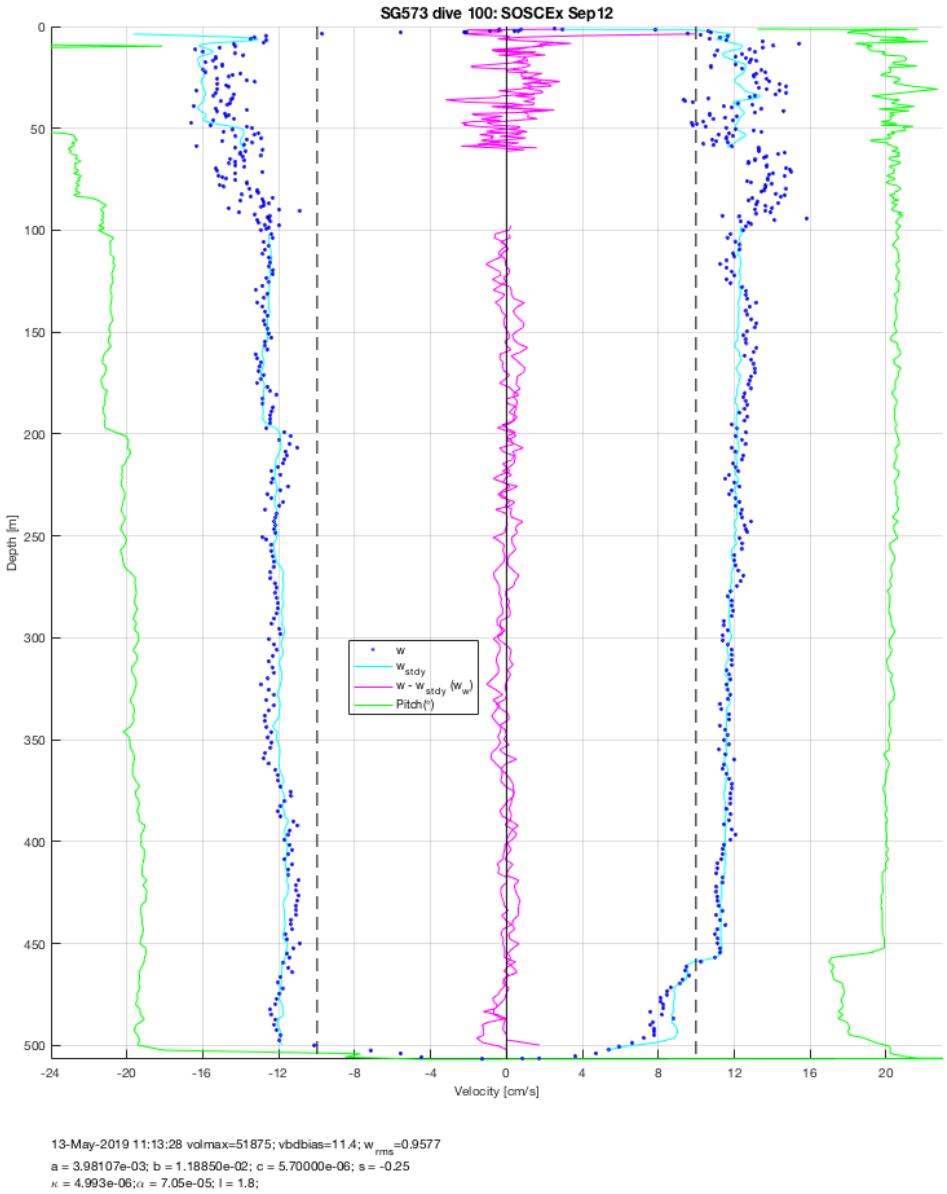

Figure 31: Vertical velocity predictions using FMS solutions for dive 100 of SG573 in the Southern Ocean in 2012.

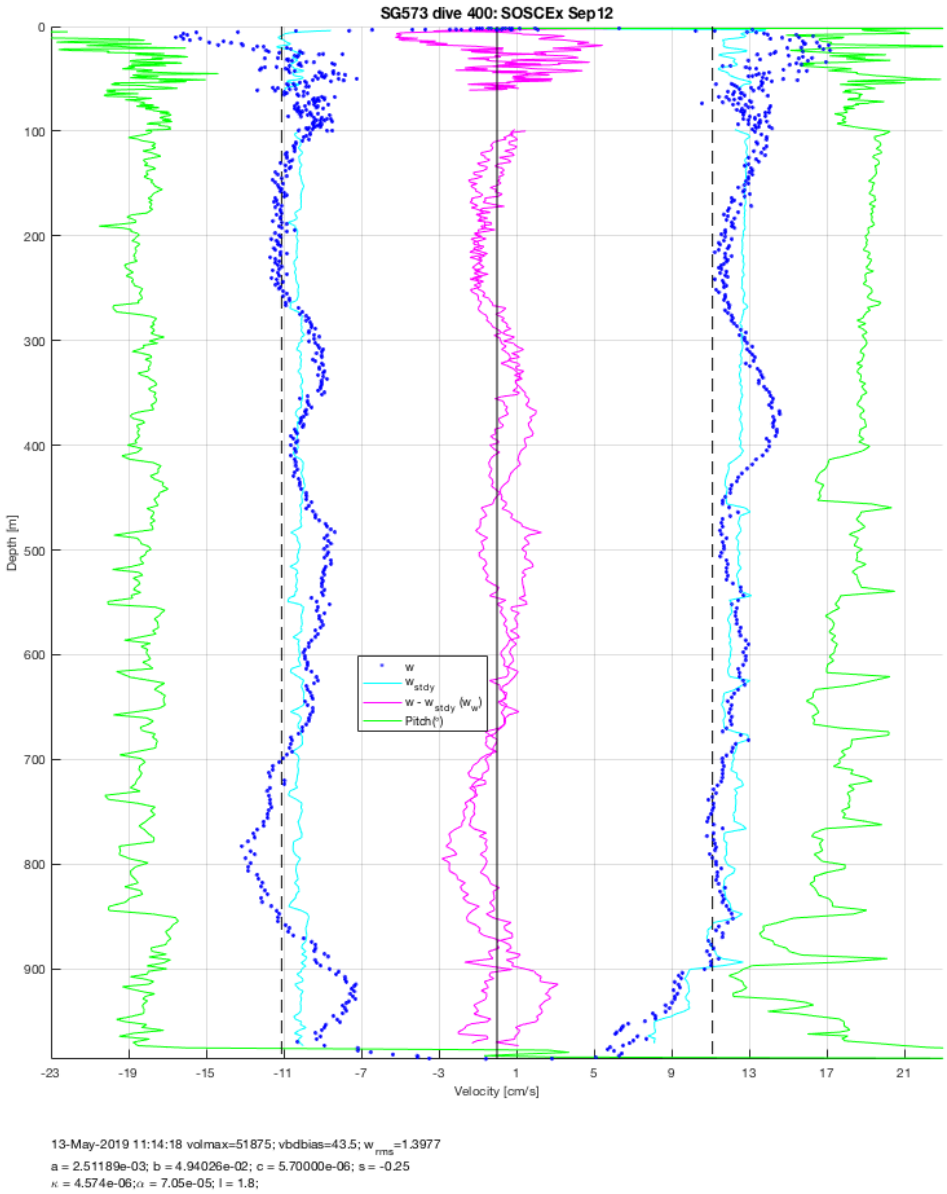

Figure 32: Vertical velocity predictions using FMS solutions for dive 400 of SG573 in the Southern Ocean in 2012. Note that drag for this dive is four times the drag of dive 100. Internal waves are visible in *w*.

CF leak on a Deepglider

During a deployment at Bermuda Atlantic Timeseries Site (BATS) in 2014, DG035 appeared to be slowly losing volume after dive 60. Post-recovery investigation showed that several of the aft containers of CF had developed small leaks; diving to pressure allowed the containers to weep fluid at roughly 1cc/day. Although this is actually a loss of M_{CF}, it appears in the regressions as a decrease in volume. Like the

small decrease in volume noted for SG144 off Ocean Station PAPA above, FMS is able to compensate for this more-rapid loss of volume. Figure 33 shows this loss of volume; Figure 34 shows no biofouling during this period, consistent with observations after recovery.

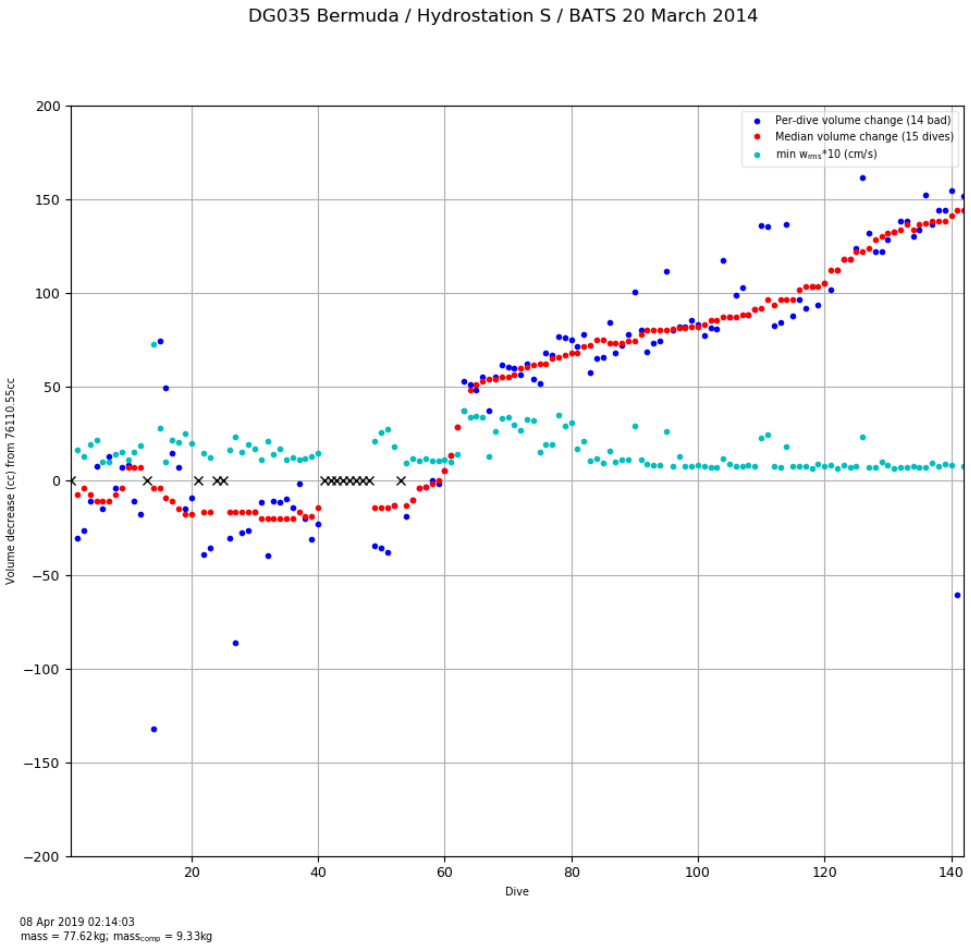

Figure 33: Per-dive solutions of δV_0 for DG035 off BATS in 2014. Blue points are per-dive estimates. Red points are a 15-dive median filtered trend. Cyan points show scaled per-dive w_{rms} values. CF loss apparently began after dive 60.

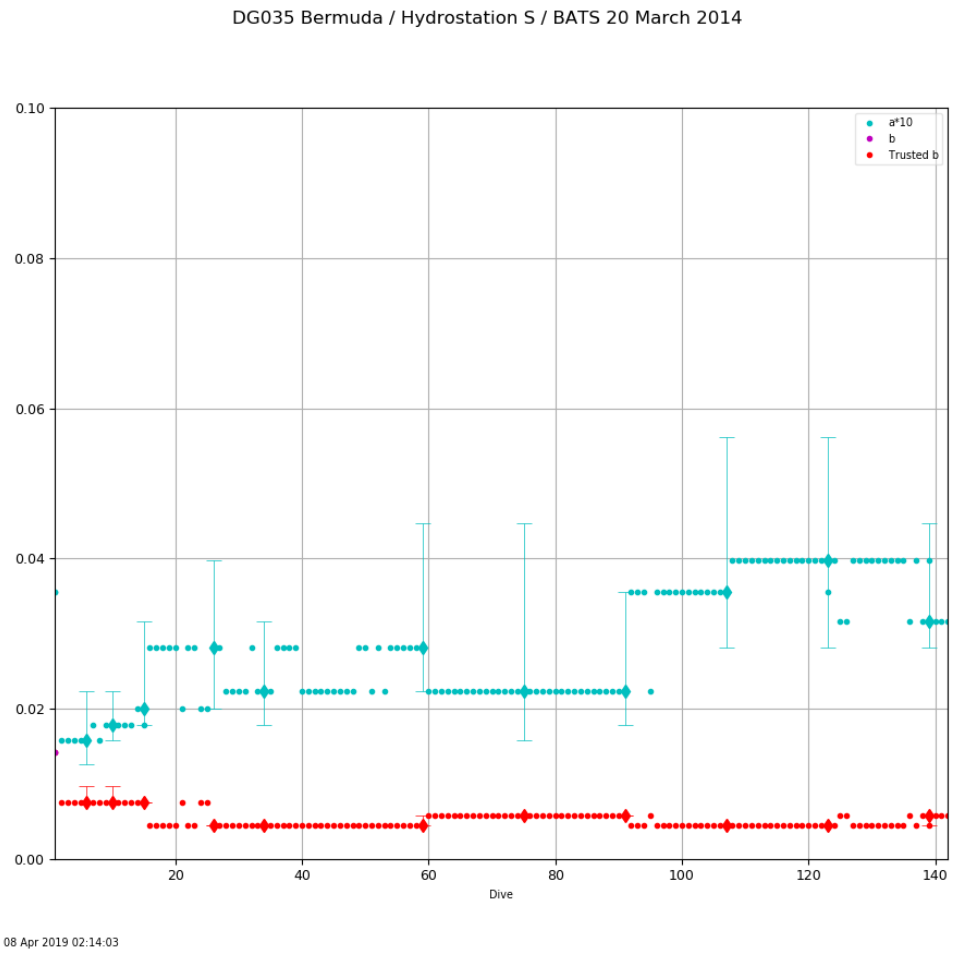

Figure 34: Per-dive lift and drag solutions for DG035 off BATS in 2014. Grid solutions are shown with diamond markers; vertical bars indicate w_{rms} solutions within 0.2 cm/s; smaller bars (and trusted markers) indicate better constrained solutions for drag. DG035 maintained a bow-tie pattern around BATS, hence the trusted solutions. There is no indication of biofouling.

Handling Antarctic Ice

As a final example of extreme volume change, consider the deployment of SG221 under the Dotson Ice Shelf in the Antarctic during 2018. Figure 35 reveals the apparent rapid volume gain of ~200cc until dive 150, when something sloughed off (losing 80cc), followed by another large and rapid loss just a few dives later. One interpretation is that the vehicle was deployed with some ice formed within or on the fairing, which increased until it broke off in pieces and stabilized after dive 200 or so. This suggests that the actual V_0 was 52682 – 150 = 52532cc. Figure 36 shows that lift and drag were unchanged except for a possible transient around dive 500. The transient might have been ice forming on the wings that subsequently sluffed off. Figure 37 shows the temperature profile vs. depth during the deployment, including the extensive under-ice excursions after dive 650.

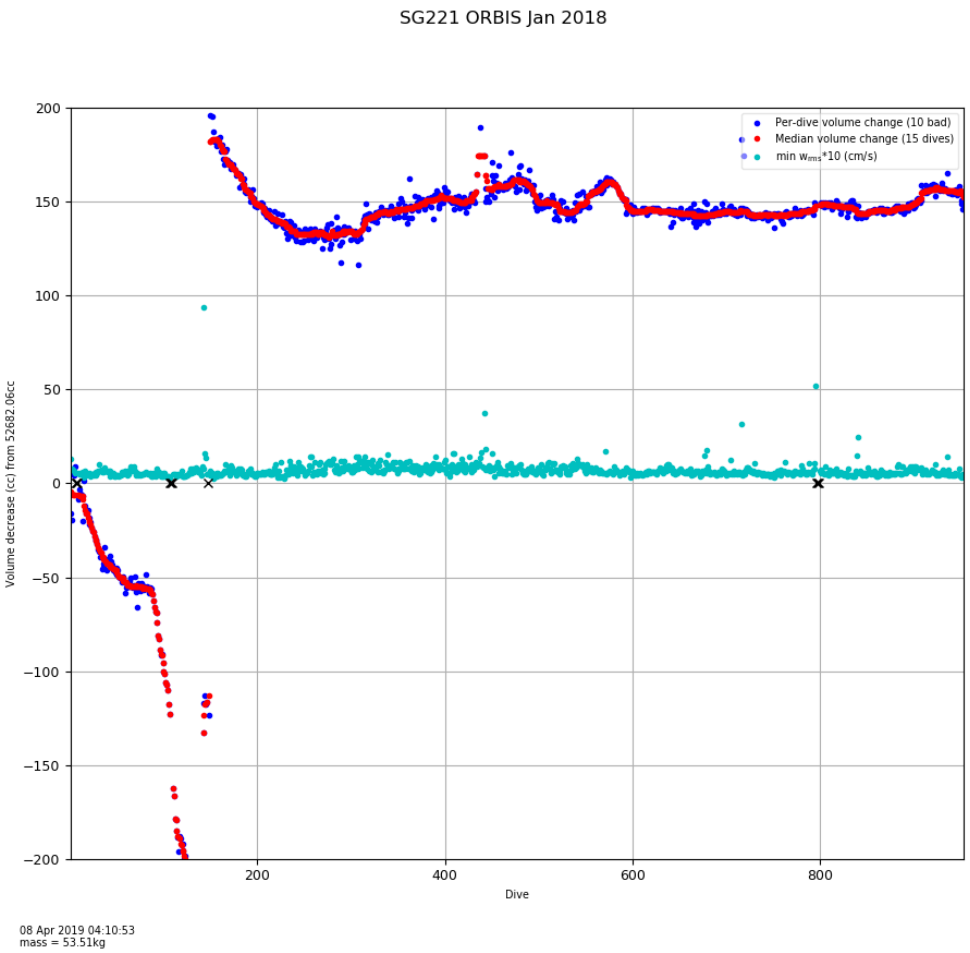

Figure 35: Per-dive solutions of δV_0 for SG221 in the Antarctic during 2018. Blue points are per-dive estimates. Red points are a 15-dive median filtered trend. Cyan points show scaled per-dive w_{rms} values.

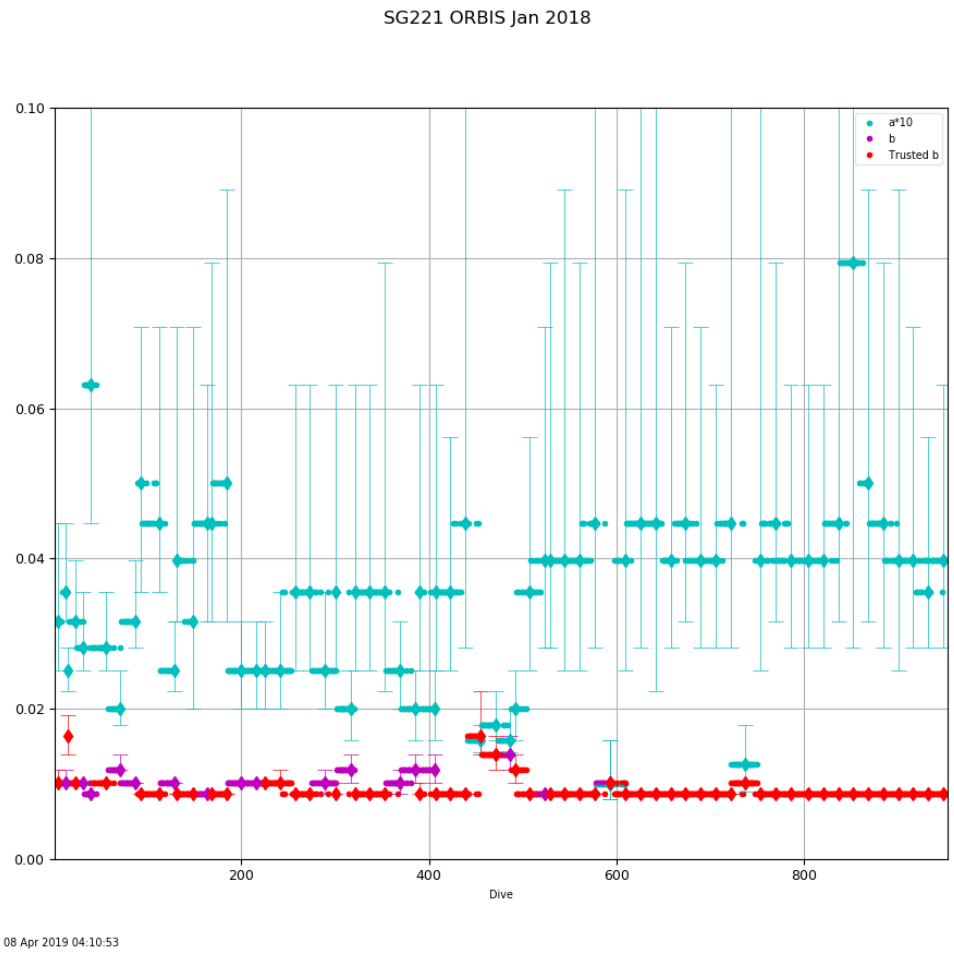

Figure 36: Per-dive lift and drag solutions for SG221 in the Antarctic during 2018. Grid solutions are shown with diamond markers; vertical bars indicate w_{rms} solutions within 0.2 cm/s; smaller bars (and trusted markers) indicate better constrained solutions for drag.

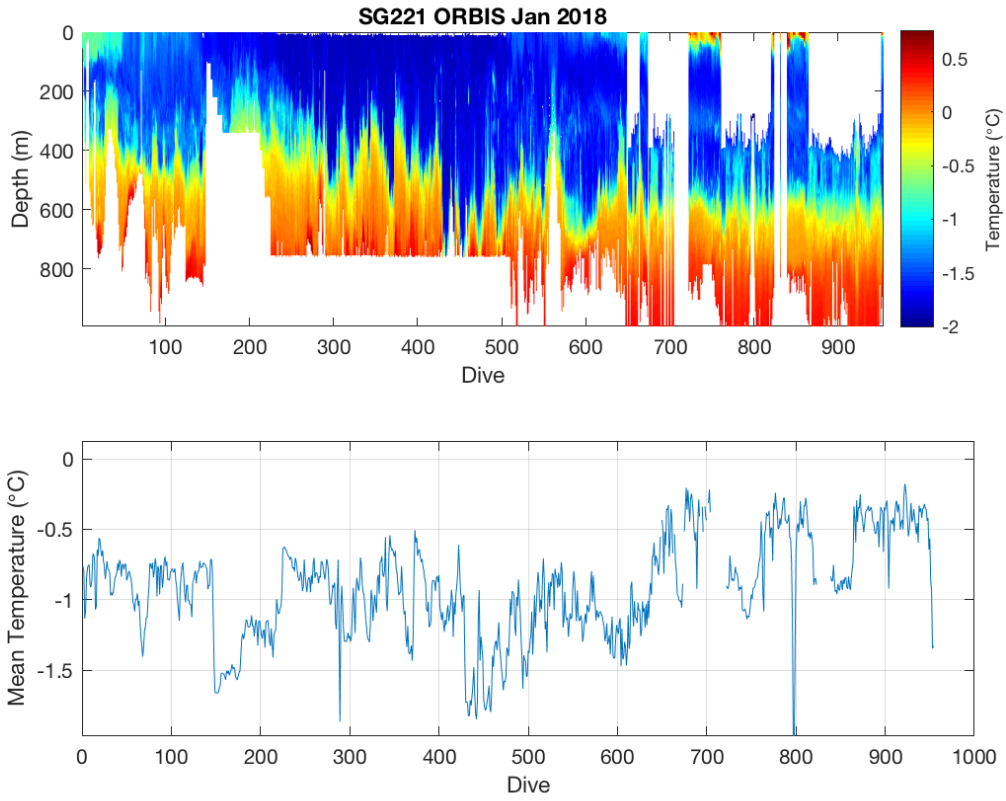

Figure 37: Seawater temperature recorded by SG221 in the Antarctic during 2018. Mean temperature is shown in the bottom panel.

Conclusions

We developed a system that automatically and incrementally determines consistent and well-constrained characterizations of the flight regimes of a glider during its deployment. Our investigations using velocimeter data suggest, however, that the *w*-only constrained solutions somewhat over-estimate the horizontal vehicle velocity. We reviewed the system's operation on several completed deployments and identified unexpected changes to some vehicles. Nevertheless, the system compensated for those issues, improving the recovered oceanographic data. There are several heuristics and thresholds that drive the system's performance. These should be investigated further through additional engineering tests and replaced with more principled solutions. However, even in its present state, consistent application of this system should improve the reliability of results and operation of glider missions.

Acknowledgments

This investigation was prompted by stimulating conversations with Kjetil Våge, Ilker Fer, Bastien Queste and the glider operations team at a glider workshop held at the University of Bergen, Norway in May, 2016; many thanks for their continued encouragement and discussion. We thank Noel Pelland for in-depth discussions of the flight model and for contributing the analysis of overall DAC impact of over/under-estimated lift and drag parameters on SG187's 2015 deployment. Craig Lee and Lou St. Laurent generously contributed additional deployment data for analysis. We thank Geoff Schilling for many comments and providing the data and explanations for operational incidents during various deployments.

References

Bennett, J. S., F. Stahr, Eriksen, C. C., 2018: An Equation of State for dodecamethylpentasiloxane, University of Washington, 16 pp.

Bennett, J. S., F. Stahr, 2014: GPS positional accuracy of Seagliders, University of Washington, 10 pp. http://hdl.handle.net/1773/43819

Eriksen, C. C., T. J. Osse, R. D. Light, T. Wen, T. W. Lehman, P. L. Sabin, J. W. Ballard, and A. M. Chiodi, 2001: Seaglider: A long-range autonomous underwater vehicle for oceanographic research. IEEE J. Oceanic Eng., 26, 424–436.

Frajka-Williams, E., C. C. Eriksen, P. B. Rhines, R. R. Harcourt, 2011: Determining Vertical Water Velocities from Seaglider. J. Atmos. Oceanic Technol., 28, 1641-1656.

Hubbard, R. M., 1980: Hydrodynamics technology for an Advanced Expendable Mobile Target (AEMT). Applied Physics Laboratory, University of Washington Tech. Rep. APL-UW 8013, 34 pp.

Merckelbach, L., D. Smeed, and G. Griffiths, 2010: Vertical velocities from underwater gliders. J. Oceanic Atmos. Technol., 27, 547–563, https://doi.org/10.1175/2009JTECHO710.1.

Merckelbach, L., A. Berger, G. Krahmann, M. Dengler, J. R. Carpenter, 2019: A Dynamic Flight Model for Slocum Gliders and Implications for Turbulence Microstructure Measurements. J. Atmos. Oceanic Technol., 36, 281-296. https://doi.org/10.1175/JTECH-D-18-0168.1

Pelland, N., C. Eriksen, and C. Lee, 2013: Subthermocline eddies over the Washington continental slope as observed by Seagliders, 2003- 09, J. Phys. Oceanogr., 43, 2025–2053. https://doi.org/10.1175/JPO-D-12-086.1

Pelland, N. A., C. C. Eriksen, and M. F. Cronin, 2016: Seaglider surveys at Ocean Station Papa: Circulation and water mass properties in a meander of the North

Pacific Current, J. Geophys. Res. Oceans, 121, 6816–6846. https://doi.org/10.1002/2016JC011920

Pelland, N. A., C. C. Eriksen, and M. F. Cronin (2017), Seaglider surveys at Ocean Station Papa: Diagnosis of upper-ocean heat and salt balances using least squares with inequality constraints, J. Geophys. Res. Oceans, 122, 5140–5168. https://doi.org/10.1002/ 2017JC012821

Rudnick, D. L., T. M. S. Johnston, and J. T. Sherman, 2013: High-frequency internal waves near the Luzon Strait observed by underwater gliders, J. Geophys. Res. Oceans, 118, 774–784. https://doi.org/10.1002/jgrc.20083

Rudnick, D.L., J. T. Sherman, A. P. Wu, 2018: Depth-Average Velocity from Spray Underwater Gliders. J. Atmos. Oceanic Technol., 35, 1665-1673. https://doi.org/10.1175/ JTECH-D-17.0200.1

Swart S., et al., 2012: Southern Ocean Seasonal Cycle Experiment 2012: Seasonal scale climate and carbon cycle links. South African Journal of Science, 108(3/4), https://doi.org/10.4102/sajs.v108i3/4.1089

Todd, R., D. Rudnick, J. Sherman, W. Owens, and L. George, 2017: Absolute velocity estimates from autonomous under- water gliders equipped with Doppler current profilers. J. Atmos. Oceanic Technol., 34, 309–330. https://doi.org/10.1175/ JTECH-D-16-0156.1

Appendix: Default flight parameter values

The terms used in Equations 1 thru 10 have direct correspondents in the basestation output (NetCDF files, *.nc) and code files. These name correspondences are listed in the table below.

M	`mass`
M_{CF}	`mass_comp`
V_0	`volmax`
δV_0	`vbdbias`
l	`glider_length`
a	`hd_a`
b	`hd_b`
c	`hd_c`
s	`hd_s`
ρ_0	`rho0`
κ	`abs_compress`
τ	`therm_expan`

T_0	temp_ref

The following measurement vectors in each dive nc file are associated with the equation terms above:

T	temperature_raw
p	pressure
ρ	seawater density computed from temperature_raw, salinity_raw, and pressure
ϕ	eng_pitchAng
z	ctd_depth
V_c	eng_vbdCC
t	ctd_time

The table below lists the initial values of flight parameters that vary between vehicle classes. Mass M and M_{CF} are provided by the pilot in sg_calib_constants.m. Other parameters are assumed fixed or estimated as described in the main paper.

Parameter	**Seaglider**	**Deepglider**	**Oculus Coastal Glider**
a	0.003548	0.003548	0.007079
b	0.011220	0.014125	0.014125
c	5.7e-6	2.5e-6	5.7e-6
s	-0.25	0	0
κ	4.10e-6	2.20e-6	2.45e-6
τ	7.05e-5	6.21e-5	7.05e-5

Appendix: Operational Considerations

FMS is integrated into Basestation release 2.12 from the University of Washington. The system can analyze dives produced under earlier versions of the Basestation but the system cannot automatically reprocess those dives with any new per-dive flight parameters; the original dive (nc) files are left unchanged in this case. However, the deployment analysis figures for δV_0, κ, and a/b shown above are produced for review.

The system maintains a database of its computations and stores various plots in a 'flight' subdirectory under the deployment directory. If the flight subdirectory is missing (or is deleted) the system will re-create the directory and (re)analyze the current set of profiles in the deployment directory.

Previously the pilot supplied values for the flight parameters in the `sg_calib_constants.m` file and was responsible for reprocessing any previous dives when those constants were updated. Under the new basestation, FMS now supplies all the flight parameters and ignores any parameters set by the pilot in `sg_calib_constants.m` file (except `mass` M and `mass_comp` M_c). When the system begins operation it first makes a backup copy of the current `sg_calib_constants.m` and then comments out any lines in `sg_calib_constants.m` that mention the flight model variables[16]. Those lines are tagged with the comment '`% FM_ignore`'. This is done to reduce confusion about where the now per-dive flight parameters originate.

Under the new basestation, FMS is triggered automatically whenever a new dive is processed during a deployment. The system can also be run separately, for example, on a completed or updated deployment; see the helper script `flight_model.sh`. In either case FMS examines all the dives currently available and determines which dives have been updated or are new to the system. These updated dives will have their flight parameters computed as needed and recorded in a flight database then applied during any reprocessing.

Initially the system applies a set of vehicle-type specific default flight parameters (see the appendix above) but as the regressions are performed updated values are applied to new dives. The system handles dives that are received out-of-order (as can happen with under-ice missions or temporary communication failures) or that perform sub-surface finishes ('yo-yo dives').

Each dive `nc` file records the specific flight parameters used for the dive and reports these values in its processing history record. If these files are reprocessed on a different basestation these flight parameter values will be retained *as long as the modified* `sg_calib_constants.m` *file accompanies the* `nc` *files* so no flight parameters are overridden by values in `sg_calib_constants.m`.

During an active deployment, FMS attempts to alert the pilot about actions that could improve the recovery of the flight parameters or the piloting of the vehicle itself. These alerts are sent to anyone electing to receive `alerts` in the `.pagers` file. In particular the system can suggest performing high-pitch-angle dives to help constrain the determination of lift and drag parameters. It will also alert the pilot to important changes to the onboard versions of the flight parameters (`$HD_A`, `$HD_B`, etc.) that are used for various types of navigation. The system alerts the pilot if initial drag is estimated to be 1.5x more than the expected default for the associated vehicle type (typically because of sensors). Finally, it can also warn the pilot about possible biofouling situations.

---

[16] Ensure that `sg_calib_constants.m` is writable by the glider account, not just pilot. Typically, this requires `chmod g+w sg_calib_constant.m`.

Various heuristic control parameters are documented at the head of the `FlightModel.py` code module in the basestation. FMS provides a simple mechanism for experimenting with these parameters. At startup FMS looks for an optional `flight_model.cnf` file first in the basestation directory and then in the deployment directory. Entries in these files provide new values, with the deployment `cnf` file taking precedence. The basestation version provides overrides for all deployments; an example is provided in the basestation release. To create a deployment-specific `cnf` file, copy the basestation version to the deployment directory and edit that copy, following the instructions in the file comments.

In particular, the policy on reprocessing dives after determining their flight parameters can be controlled in this manner. FMS is relatively quick at making its (incremental) estimates of flight parameters. However, the reprocessing of dives using the updated parameters can be expensive in processing time, especially if several dives need to be reprocessed at once. The reprocessing can be avoided (on a basestation-wide or deployment-specific basis) by uncommenting the `enable_reprocessing_dives: False` line in the appropriate `cnf` file. The system will stop reprocessing but will continue to estimate per-dive parameters and store them in the flight subdirectory until the line is again commented out or the value is changed to `True`. Dives received after disabling the system will use the flight parameters of the most-recent dive as usual. When the system is re-enabled, all dives with changed parameters will be reprocessed when FMS next runs.