

Subject: Recommendation re calculating angle-of-attack

Al Cooper

16 November 2016

Background

Since NCAR/RAF started using ports in radomes to measure angle of attack and sideslip, the relationship has been represented by empirical equations that relate observed variables like ADIFR to the desired variables like AKRD. Originally, these relationships were based on expectations for hemispherical geometry, and early radomes were approximately that shape, but this was found to be impractical for later aircraft and especially for high-speed aircraft like the Sabreliner and GV, which have radomes designed to minimize drag that depart significantly from hemispherical shape. The approach to “calibration” of radomes now relies on collecting measurements where the vertical wind is thought to be small and finding empirical relationships that lead to near-zero vertical wind for those regions. Particularly useful are speed runs because the angle of attack varies over a wide range while the maneuver is flown maintaining altitude in a region of low vertical wind. Fitting to data from such maneuvers can lead to a good representation of the angle of attack if the fit is required to give near-zero vertical wind. However, it has often been difficult to include speed runs that cover the necessary range of conditions during research flights for a particular project, so it has often been necessary to use larger data sets spanning many flights and conditions in order to seek to represent the vertical wind in cases where historical relationships have appeared to provide poor results.

The current approach

For recent projects, the following representation has been used, as described in the document on processing algorithms and in the NCAR Tech Note on Wind Uncertainty:

$$\alpha = c_0 + \frac{\Delta p_\alpha}{q} (c_1 + c_2 M) \quad (1)$$

where α is the angle of attack, Δp_α is the pressure difference between top and bottom ports on the radome (i.e., ADIFR), q is the dynamic pressure (QCF has been used for reasons explained in the Processing Algorithms document), and M is the Mach number calculated from dynamic and static pressure before static-defect corrections and without inclusion of effects of humidity. In a region with negligible vertical wind, the reference value α^* (AOAREF, expressed in degrees) used in fitting to find the coefficients $\{c_i\}$ is¹

$$\alpha^* = \theta - \frac{180}{\pi} \arcsin \frac{w_p}{V} \quad (2)$$

¹ $w \approx V \sin(\alpha - \Theta) + w_p$ so, setting $w = 0$ leads to this formula.

where θ is pitch (PITCH), w_p is the rate-of-climb of the aircraft (GGVSPD), and V is the airspeed (TASX). The fit coefficients are normally determined in units of degrees to give α (AKRD) in degrees, so the arcsin function is multiplied by $(180 / \pi)$ to convert that angle to degrees also.²

During the reprocessing effort, it was found that the empirical coefficients need to provide reasonable vertical wind often varied from project to project and sometimes within a project, and this was puzzling because, although the aircraft configuration and loading varied among projects, it did not seem likely that the radome performance changed so much. There was a radome replacement in January 2013, prior to the SAANGRIA-TEST project, so it is reasonable that earlier check- projects appear to need significantly different coefficients. However, difficulties with vertical wind have persisted through some recent projects including CSET and ORCAS, so it appears that some unknown problem is causing the representation of angle of attack to change.

Reason for proposing a change

Consistent results for most projects, except when the radome changes, seem desirable, so there appears to be some weakness in the current approach. There are occasional problems with the radome measurements themselves, such as leaks or plugged lines, but those are usually obvious and such measurements are flagged as bad in data review, so those effects are not likely related to this problem in most projects. Instead, the most likely origin of the problem may be an offset in flow angle at the radome created by some aerodynamic effect of the aircraft configuration (esp. wing pods) or flight-control factors like the amount of lift generated by the wings, thrust generated by the engines, or changes in control surfaces (flaps, speed brakes, trim settings).

The suggestion investigated in this memo is that some significant part of the problem arises from the need to represent two different components in AKRD with one formula. Theoretical considerations suggest that the response to gusts (here, short-period fluctuations) should be consistent and described well by a formula like (1) with $c_2 = 0$. However, our difficulties have been instead with the longer-period response that leads to an apparent offset in the zero reference point for AKRD and therefore a bias in vertical wind. These slower variations may originate in aerodynamic effects that vary among projects and among flights and so are difficult to represent well with an equation like (1). That equation was originally introduced to improve the range of altitudes over which the equation provided reasonable fits and especially to correct for an apparent offset at low altitude, and the introduction of this particular representation involving M was chosen after investigating many other possibilities. However, representing both short-period fluctuations and the zero reference value with one formula causes the resulting representation to be compromised in how it represents both. Evidence to be presented below suggests that the gusts are represented well with a consistent formula that applies universally, but that the problem arises from efforts to maintain a good zero reference point and that latter problem requires a different empirical equation.

²In fits, $\arcsin(w_p/V)$ has often been replaced by w_p/V because the difference for arguments of a few degrees is negligible.

One approach, used in some past projects like HIPPO, has been to high-pass-filter the measurements of vertical wind to remove the slowly drifting offset. This makes it impossible to detect real long-period non-zero values of vertical wind but leaves the shorter period fluctuations intact. Most research relying on the vertical wind measurements has not needed long-term accuracy but rather accurate measurement of gusts, so filtering is a possible but certainly unsatisfying compromise. A different approach, developed in this memo, is to split the measurements into two components by complementary filtering into high-pass and low-pass components and then treat these two separately. The two components are then added back together to provide a measurement of AKRD that is not compromised by filtering and can detect slowly varying components of AKRD and vertical wind.

The suggested change

The reference value α^* , which is always valid at all scales and in all conditions if the vertical wind is indeed negligible, is split into two components $\alpha^* = \alpha_f^* + \alpha_s^*$ that result from applying a Butterworth low-pass filter to α^* to obtain α_s^* and then finding α_f^* from $\alpha_f^* = \alpha^* - \alpha_s^*$. These components are then represented by separate fits:

$$\alpha_f = c_1 \left(\frac{\{\text{ADIFR}\}}{\{\text{QCF}\}} \right)_f \quad (3)$$

$$\alpha_s = d_0 + d_1 \left(\frac{\{\text{ADIFR}\}}{\{\text{QCF}\}} \right)_s + d_2 \{\text{QCF}\}_s \quad (4)$$

where the f and s subscripts represent the high-pass and low-pass components after filtering. It should be the case that a constant term is not needed in (3) because any offset should be removed by filtering, and that was verified by including a constant c_0 in the fit and finding it to be negligibly small. The term QCF was introduced after finding that many of the apparent errors in vertical wind occurred when the estimated lift (proportional to dynamic pressure) was different from that expected for normal flight. Introduction of that term led to significant improvement in the representation of the slow component, but not the fast component. More complicated representations were tested in both cases, but these appear to provide adequate fits without additional terms, as shown below.

This approach has two advantages:

1. The important sensitivity to fluctuations is not compromised by efforts to represent the slowly varying zero level for angle of attack.
2. The slowly varying zero reference can be represented by a more complex equations without needing to apply those to the high-frequency component. Separating this component allows much better representation of the offset in AKRD and also includes valid representation of a difference from the offset arising from a true low-frequency component of vertical wind.

The coefficients in (4) can be determined by fits to several projects and many flights, so an individual low-frequency vertical wind will still be detectable as a departure from the fit constrained by that large data set.

Details of the fits

Some aerodynamic evidence

The results for vertical wind in ORCAS were particularly problematic, and many of the problems appeared associated with periods during or shortly before climbs and descents when changes in control of the aircraft were likely. (See the memo AKRDforORCAS.pdf.) To investigate the effect of flight conditions, the following term K was introduced that should characterize the performance and might correlate with the observed errors. It was calculated as follows:

1. The lift L generated by an airfoil is expected to depend on a coefficient of lift c_L , the air density ρ_a , the airspeed V , and a representative surface area S , according to:

$$L = \frac{1}{2} \rho_a V^2 c_L S \quad (5)$$

and in turn the coefficient of lift depends approximately linearly on the angle of attack α , although this angle of attack may be offset from the standard value relative to the aircraft longitudinal axis.

2. The relationship and reference angle for angle-of-attack may change when flaps are deployed so it might be expected that $c_L = c_0(\alpha - \alpha_0)$ where c_0 and α_0 both change with flap deployment. The factor $c_L S$ therefore changes with changes in aircraft configuration. The lift will normally be matched approximately to the weight of the aircraft W except during turns and short periods of vertical acceleration at the start and end of climbs and descents.³ In level or constant-vertical-speed turns, the lift needs to be increased to $W/\cos\phi$ where ϕ is the roll angle. Then from (5) it would be expected that

$$c_L S = \frac{2W(t)}{\rho_a V^2 \cos\phi} = \frac{W(t)}{q \cos\phi} \quad (6)$$

where $q = \rho_a V^2/2$ is the dynamic pressure. K is defined to be

$$K = \frac{A_k W(t)}{q \cos\phi} \quad (7)$$

where the factor $A_k = (500 \text{ lb})/W_0$ is introduced for plotting convenience.

³A first-order correction would be $L(1 + \gamma^2) = W$ where $\gamma = \arctan(w_p/V)$ with w_p the rate of climb and V the airspeed.

3. The aircraft weight $W(t)$, needed to calculate K , can be estimated from $W(t) = W_L - B_f(t - t_L)$ where W_L is the weight at landing time t_L and B_f is the fuel burn rate, typically around 1 lb/s for the GV.⁴
4. For normal flight (without flaps or speed brakes) S would be constant and the coefficient of lift would vary linearly with α , according to $c_L = c_0 + c_1\alpha$, leading to a region in a plot of K vs. α with a linear relationship. For normal flight, a scatterplot of these variables (cf. AKRDforORCAS.pdf) showed that they can be regarded as linearly related without offset, approximately according to the relationship $K \approx 1.5\alpha$.
5. For “normal” flight this leads to the expectation that

$$c_1 S \propto \frac{K}{\alpha}$$

(with the \propto sign indicating proportionality) and therefore to the expectation that K/α will be constant unless some factor like deployment of flaps caused a change in the coefficient of lift. Determining α is the problem being addressed, but α^* can be used as a surrogate with the assumption that the vertical wind is small. In that case, K/α^* can be used as an indicator of flight conditions, with an expected value of about 1.5.

Direct use of these equations would be difficult for routine calculation because it would be necessary to know the flight-dependent aircraft weight and its variation with time. However, an example can illustrate the usefulness of the variable K , as obtained from ORCAS flight 6 by assuming a landing weight of 65,000 lb and a fuel burn rate of 0.85 lb/s. For that case, Fig. 1 shows K/α^* and shows that, during the descent, low-level-leg, and climb from 20:00 to 21:00 UTC there is a departure from the expected normal-flight value that is significant. There are also departures during the initial climb and final descent, with brief periods where the value substantially exceeds 1.5, which are likely periods where flaps and/or landing gear were deployed.

Incorporation of the factor K into empirical fits indeed led to significant improvement in the fits, but this has not been pursued further because of the difficulty associated with calculating K . Instead, this has been taken to be an indication that some aspect of aerodynamic effects influences the vertical wind and needs to be represented. K depends inversely on q , so one indication arising from this result is that some function of q (other than ADIFR/QCF) might improve the fit results, especially for the slowly varying component of angle of attack.

⁴A better estimate would consider variations in this fuel burn rate during climbs and descents and with aircraft weight, but this is a reasonable first estimate. The python program "PlanFlight.py" incorporates better estimates of the rate of fuel burn.




Figure 1: The aerodynamic factor K divided by the reference value of angle of attack, for OR-CAS flight 6.

Filtering

Two approaches to filtering were tested in order to determine if the difference was significant. Both were similar Butterworth low-pass filters applied to α^* and also to the factors used to determine the slowly varying component of angle of attack, but one combined forward and backward filtering in order to minimize phase shifts at the transition frequency while the other used only forward filtering. The reason for testing both is that the former is a preferable method but the latter can be incorporated into first-pass processing in nimbus without the requirement for a second pass to accommodate the filtering. The specific filter used was the Butterworth filter provided by the “signal” package in R via the “filtfilt” function (using “butter” within that filter function), with a cutoff frequency of 1/600 Hz. This cutoff frequency was selected because it typically corresponds to around 100 km wavelength and so likely smooths out most real turbulent fluctuations while leaving a long-term mean value. Independent variables entering the fit, including ADIFR/QCF, QCF, Mach Number, and others used in test runs, were also filtered. The “fast” component of each was then determined by subtracting the filtered variable from the unfiltered variable.




Figure 2: [top] Reference value for angle of attack, showing single-pass low-pass filtered result (S1), two-pass low-pass filtered result (S2), and the high-pass result for both cases (F1, F2), for ORCAS flight 6.

[bottom] The sum of the low-pass and high-pass components for the single and double pass filters, showing that the result is the same for both as expected.

Figure 2 shows the result of filtering in this way, either single-pass or forward and backward. There is a substantial difference between the results filtering single vs two-pass, but of course no difference once the components are re-combined. Representing the slow component by fits determined from data from a single flight may be ineffective, but it could still be the case that such fits using composite data from many projects and flights could average in ways that provide a reasonable representation. Alternately, the fit coefficients determined from double-pass filtering applied to composite data sets could be used with single-pass processing, if evaluation indicates that results this way are not distorted significantly from double-pass processing results. Therefore it is useful to evaluate the effects of these different processing schemes. That will follow discussion of the fitting procedure itself.

Finding the coefficients

To find appropriate coefficients to use with (3) and (4), a composite data.frame was constructed using flights from ORCAS, CSET and DEEPWAVE. Some selection of flights was made to exclude ones where there appeared to be possible problems, to avoid distortion of the fit provided by the best flights. For CSET, the selected flights were 1, 2, 3, 4, 6, 10, 12, and 16. For ORCAS, flights 1, 2, 3, 6, 8, 9, 11, 14, and 16 were selected. For DEEPWAVE, all 26 flights were used excepting 6, 7, 15, and 23.

Each flight was processed individually, including definition of new variables like $QR=ADIFR/QCF$ and $AOAREF$ and filtering to separate the low-frequency and high-frequency components of each variable used in the fits. Because the filter introduced start-up perturbations, periods within 600 s of the start and end of each flight were eliminated from the data.frame. Also, the results were limited to $TASX > 110$ and $abs(ROLL) < 2$, to avoid periods in turns and near the start and end of flights where use of flaps might distort the results. The result was a data.frame containing more than 800,000 measurements.

The high-pass component

For the high-pass component, Fig. 3 compares the result of the fit to that of the reference value. The residual standard deviation for this fit was 0.113° , but this includes fluctuations caused by real vertical wind (corresponding approximately to a standard deviation in the vertical wind of 0.4 m/s), so this is not unexpected for a fit to such a large data set. There are so many points that a scatterplot such as that shown just becomes a blob that is difficult to interpret in terms of probability, so the field in this plot was binned into a 100x100 array and contours were drawn for that array, as shown by the red lines. They show that most of the measurements are confined to small departures from zero in both variables and that those small values will have a strong effect on the fitted slope.

A previous study for DEEPWAVE, discussed in the Tech Note on Wind Uncertainty, led to a slightly different value: $c_1 = 21.481^\circ$. This also gives a good representation of the measurements, with the residual standard deviation increased by only 0.0005° over that for the best fit. This result was based on speed runs and other carefully qualified measurements from that project, so it may

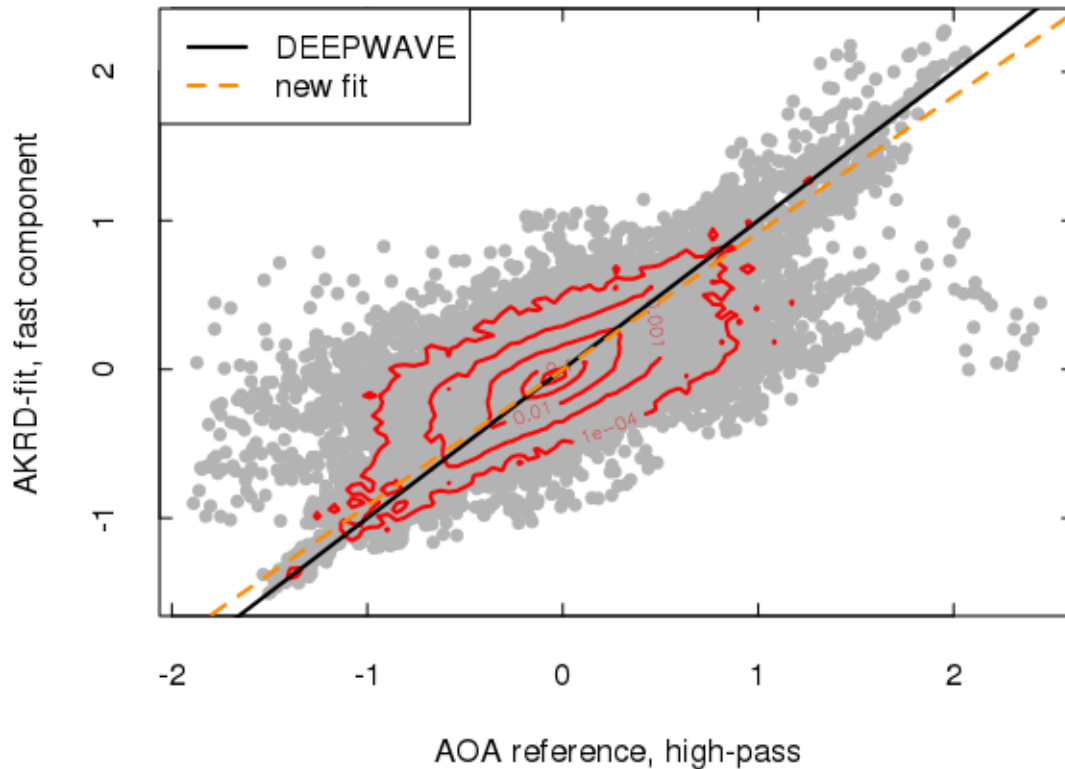


Figure 3: For the high-pass component, the fit result as a function of the reference value, for all points entering the fit. The contour lines denote probability levels relative to the maximum bin, which was assigned the value of 1. Contour lines were drawn at these values: 0.5, 0.1, 0.01, 0.001, 0.0001.

be preferable to use that value. The black line in Fig. 3 represents the DEEPWAVE fit, while the dashed orange line is the fit determined from the comprehensive data set used in the present study. The DEEPWAVE value appears to represent the values at large and small angle of attack better than the orange-line fit, so the DEEPWAVE value will be used in the rest of this study.

The low-pass component

The fit to the low-pass component led to significantly lower residual standard deviation, 0.059° , than for the high-pass component, presumably because the latter is affected by real non-zero vertical wind while the former is not. A scatterplot like that shown for the high-pass component shows little scatter, because the fit accounted for 98.3% of the variance. Instead, a plot of the fit result

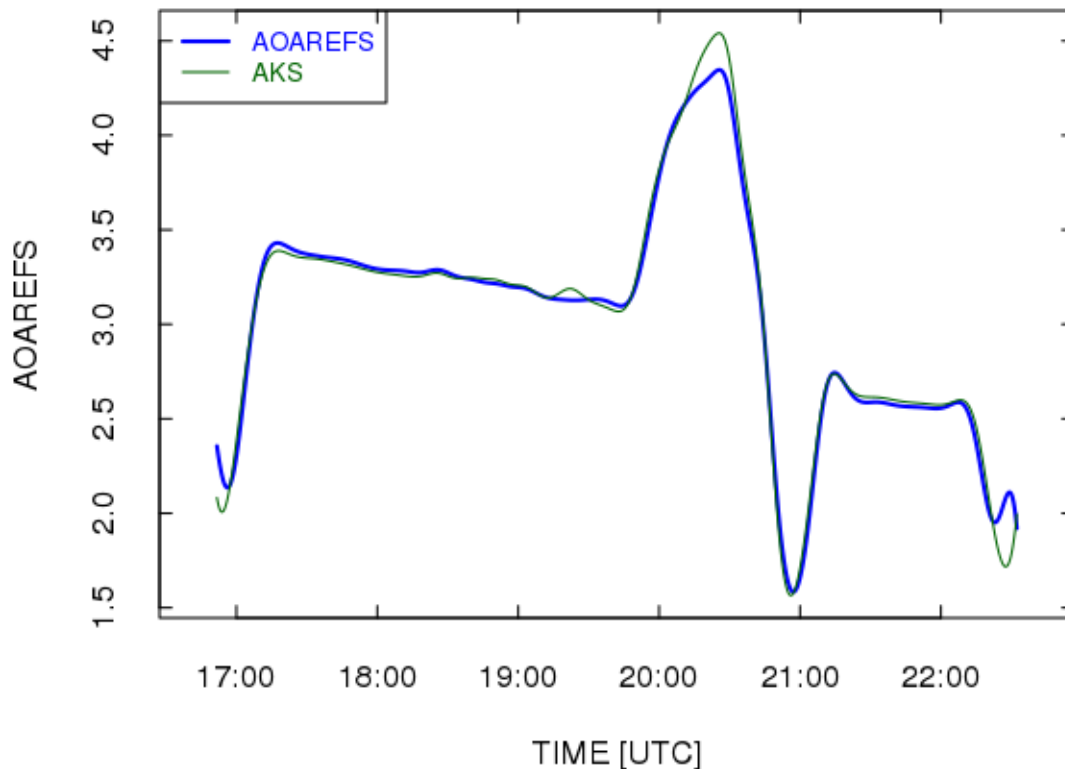


Figure 4: The calculated low-pass component of angle of attack (AKS) and the reference value used for the fit (AOAREFS), for ORCAS flight 6.

compared to the reference result for individual flights is more informative. Figure 4 shows such a comparison for ORCAS flight 6. There is a significant perturbation near 20:20 UTC, which may be associated with the extended turns that take place at this time during a racetrack descent or with other aerodynamic effects, but otherwise the representation of the slow component is very good. Other flights from this and other projects look similarly good.

Summary

The resulting representation of AKRD is

$$\{AKRD\} = c_1\{QRF\} + d_0 + d_1\{QRS\} + d_2\{QCFS\} \quad (8)$$

where QRS is the low-pass-filtered ADIFR/QCF, QRF is the high-pass version obtained by subtracting the low-pass version from the original variable, and QCFS is the low-pass-filtered QCF.

The coefficients are $c_1 = 21.481$ and $\{d_i\} = \{4.5253, 19.9332, -0.00196\}$. The coefficients c_1 and d_1 are close to the same value, so the result is that the factor ADIFR/QCF is applied in about the same way to both components of AKRD. The advantage of the complementary-filter approach is that the term QCF is applied only to the slow component to adjust the offset in AKRD while not affecting the high-frequency response.

Results for the vertical wind

The second-pass approach

To use (8) in the calculation of vertical wind, the following steps are followed:

1. Construct a list of all files for which addition of a new vertical-wind variable or substitution for the existing variable is desired. Multiple projects can be included in the list, and the available files for a project can be added to the list using the R function “list.files()” as in the code attached to this memo, which calculates new values for the MPEX project.
2. Loop over each file in that list:
 - (a) Copy the file to a new file with “F” appended to the name; e.g., MPEXrf01F.nc, so that any work on the netCDF file does not affect the original. Make all changes to the new file.
 - (b) Check that the variables needed for processing are present in the netCDF file, and select an appropriate variable to represent the rate-of-climb of the aircraft. The variables needed include all those on which the three-dimensional calculation of wind depends, because the normal calculation will be repeated when constructing the new value for vertical wind. If any is missing, skip the calculation for that file. For rate-of-climb, recent projects use GGVSPD but some have used GGVSPDB or VSPD_A or even VSPD_G so the appropriate variable needs to be selected for the calculation. In straight flight without roll, a good approximation would be to calculate the new vertical wind from $WIY = WIC + (AK - AKRD) * TASX * \pi / 180$, but a correction is needed in turns where the non-zero roll is important in the calculation, so the full three-dimensional calculation is warranted. In addition, this is an opportunity to add the effect of the distance that separates the GPS antenna from the INS, which has not been included in the processing leading to the files on /scr/raf_data, where the sources for this calculation reside.
 - (c) Open the netCDF file with the “write” option, and from the dimensions of the “Time” variable determine the data rate. This is needed during the filtering step to maintain the same effective transition frequency for high-rate and low-rate files.
 - (d) Define new variables $QR = ADIFR/QCF$ and $AOAREF = PITCH - (180/\pi) * GGVSPD / TASX$, or an appropriate substitute for GGVSPD where required. For these variables and also for QCF, find the low-pass-filtered version using a cutoff frequency of

- (1/600) Hz. For QR and AOAREF, also find the high-pass version by subtracting the low-pass-filtered variable from the unfiltered variable.
- (e) Find a new angle of attack (AK) using (8).
 - (f) Use the Ranadu function “WindProcessor()” with the resulting angle-of-attack to find a new vertical-wind variable, here called WIY. Also apply the same filter to WIY to obtain the high-pass version WIF, for inclusion in the files if desired.
 - (g) Add the new variables to the netCDF file, with appropriate attributes. At present, the new angle of attack is also saved as the variable AKY. If the original file is high-rate (25 or 50 Hz), add the new variables at this rate also.
3. Construct a memo documenting the process. This can be done using the file AKRD.Rnw, which does the processing described above while constructing a documenting memo

The result is a new file containing the new vertical-wind variable WIY and the high-pass-filtered version WIF, along with the new angle of attack AKY. Once this is reviewed, if accepted for processing I think it would be better to produce the variable WIC in this way. The name WIC is awkward now: Attributes say “GPS-corrected” but it is not the result of a correction to another variable but rather a variable based wholly on GPS. The “C” would be a natural indicator of the complementary-filter process. Other variables might provide alternatives for vertical wind, like the “WIR” variable proposed separately. It might then be useful to introduce a variable WIX that denotes the preferred variable, as has been done for HIPPO flights. AKRD should eventually be based on this approach also.

Processing to produce this new variable has been performed for these projects: ORCAS, CSET, DEEPWAVE, MPEX, and CONTRAST, and the results look very good for all these projects while using the same coefficients listed above. The procedure does not work for DC3 or TORERO, I presume because there has been a radome change or other change after those projects that caused an offset not present in recent projects. It is a valuable development, however, that one processing approach works this well for all these recent projects, some of which have had serious problems with the calculation of vertical wind.

The single-pass approach

The approach developed above depends on post-processing after nimbus has run. It is worth exploring if this might be incorporated into nimbus through use of a recursive filter instead of the two-pass filter used above. It was shown in Fig. 2 that using the same Butterworth filter in a recursive version leads to a significant phase shift in the low-pass component and the introduction of corresponding differences in the high-pass component. However, after combining the low-pass and high-pass components of the resulting AKRD, the variable might still be valid. To test this, single-pass filtering was used to produce test files that were compared to the forward-and-backward-filtered version. In the ORCAS files [ORCASrffxF.nc], there is also a variable WIS that was calculated with a recursive filter. A comparison of WIS to WIY, obtained with a forward-and-backward filter,




Figure 5: A comparison of results obtained by forward-and-backward filtering (WIY) vs single-pass recursive filtering (WIS), for ORCAS flight 6. Also shown is the difference $10*(WIS-WIY)$, as the red line with the factor of 10 applied to show the magnitude of the difference.

is shown in Fig. 5. The two results are barely distinguishable on this plot, although the difference (magnified by a factor of 10 and shown by the red dashed line) shows that there is sometimes a difference of 0.1 m/s. However, it is slowly varying and does not affect the high-frequency signal to any significant degree. The respective mean values for WIY, WIS, and $WIS-WIY$ are -0.03 , -0.03 , and 0.0003 m/s, with full-flight standard deviations of 0.33, 0.33, and 0.06 m/s. Results for other flights are similar, so the single-pass result should be fully acceptable and avoids the need for second-pass processing.

Recommendation

Implement the single-pass processing described above as standard in nimbus. These changes will be needed:

1. Implement a filter similar to that now used for the horizontal-wind complementary filter, but with cutoff frequency at (1/600) Hz. As implemented before, the filter coefficients generated for a Butterworth filter by the R routine `signal::filter` were transferred to nimbus, I think by Janine. It was important to transfer them at full double-precision because the operation of the filter was found to be very sensitive to the accuracy of these coefficients. The low-pass filter will need to be applied to (ADIFR/QCF) and to QCF. The high-pass result for (ADIFR/QCF) is also needed and can be obtained by subtracting the low-pass result from the unfiltered variable.
2. Calculate angle of attack using (8) with coefficients as listed in the summary section above, Sect. . This formula must be applied to the results for the filtered components, the low-pass components for d_i and the high-pass component for c_1 .
3. Replace WIC with this new variable, and replace AKRD with the new result for angle of attack.
4. Optionally include a variable WIF which is the high-pass-filtered version of WIC.
5. Optionally include a new variable WIX which is assigned to the recommended variable for vertical wind, usually WIC. (But see a separate memo regarding introduction of a backup variable, WIR, based on rate-of-climb determined by integration of the hydrostatic equation with updating to the GPS velocity. This might be an alternative variable if there are ever problems with the GPS-based rate of climb.)

Those reviewing this memo can find reprocessed files in directories named “WAC” under the standard project directories in /scr/raf_data. There are new variables WIY and WIF representing the results of processing for the projects ORCAS, CSET, CONTRAST, DEEPWAVE, and MPEX. For ORCAS only, there is also a variable WIS representing what would be calculated using the single-pass process recommended above.

The routine that generated all these files, named AKRD.Rnw, is on EOL file space under ~coop-erw/RStudio/Reprocessing. If desirable, that could be modified to replace WIC and AKRD in the existing netCDF files, as an alternative to nimbus reprocessing as recommended above. That routine also contains an outline of how to update the calibration coefficients if that should become necessary.