

Subject: Complementary-filter approach to aircraft motion

Al Cooper

9 October 2011

Background

Wind measurements combine a measurement of relative wind with a measurement of aircraft motion to determine the air motion relative to the ground. The aircraft motion has long been measured by an IRS, and recently also by a GPS. These have complementary strengths: The IRS provides very good information on short-term motion but drifts with a characteristic period of more than an hour, while the GPS provides good absolute accuracy but sometimes is unable to receive the GPS signals and (except in differential-GPS mode) can have short-term errors that make short segments of the track look jagged. To take advantage of the strengths of each, a complementary-filter calculation was developed and implemented in the 1980s, but it was never documented publicly. This memo is partly an attempt to remedy that and partly a suggestion to make some minor changes to how it is implemented.

Aircraft velocity

To accomplish this, a low-pass filter, $F_L(\{GVNS, GVEW\})$, is applied to the GPS measurements of groundspeed, $\{GVNS, GVEW\}$, which are assumed to be valid for frequencies at or lower than the cutoff frequency f_c of the filter. Then the complementary high-pass filter, denoted $(1 - F_L)(\{VNS, VEW\})$, is applied to the IRS measurements of groundspeed, $\{VNS, VEW\}$, which are assumed valid for frequencies at or higher than f_c . Ideally, the transition frequency would be selected where the GPS errors (increasing with frequency) are equal to the IRS errors (decreasing with frequency). The filter used is a three-pole Butterworth lowpass filter, coded following the algorithm described in Bosis, S. M., 1980: *Digital and Kalman filtering : An Introduction to Discrete-Time Filtering and Optimum Linear Estimation*, p. 49. The digital filter used is recursive, not centered, to permit calculation during a single pass through the data. If the cutoff frequency lies where both the GPS and INS measurements are valid and are almost the same, then the detailed characteristics of the filter in the transition region (e.g., phase shift) do not matter because the complementary filters have cancelling effects when applied to the same signal. The transition frequency f_c was chosen to be (1/600) Hz (but this value can be overridden via the “defaults” file). The Butterworth filter was chosen because it provides flat response away from the transition. The resulting variables for aircraft motion, $\{VNCS, VEWC\}$, are then each the sum of two filtered signals, calculated as described in the following box:

VEW = IRS-measured east component of the aircraft ground speed
 VNS = IRS-measured north component of the aircraft ground speed
 GVEW = GPS-measured east component of the aircraft ground speed
 GVNS = GPS-measured north component of the aircraft ground speed
 $F_L()$ = three-pole Butterworth lowpass recursive digital filter

$$\{\text{VNSC}\} = F_L(\{\text{GVNS}\}) + (1 - F_L)(\{\text{VNS}\})$$

$$\{\text{VEWC}\} = F_L(\{\text{GVEW}\}) + (1 - F_L)(\{\text{VEW}\})$$

This is straightforward and effective when both sets of measurements (IRS and GPS) are available. The approach in use becomes more complicated when the GPS signals are lost, as sometimes happens in sharp turns. Then some means is needed to avoid sudden discontinuities in velocity (and hence windspeed), which would introduce spurious effects into variance spectra and other properties dependent on a continuously valid measurement of wind. To extrapolate measurements through periods when the GPS measurements are not available, a fit is determined to the difference between the best-estimate variables $\{\text{VNSC}, \text{VEWC}\}$ and the IRS variables $\{\text{VNS}, \text{VEW}\}$ for the period before GPS reception was lost, and that fit is used to extrapolate through periods when GPS reception is not available. The procedure is as follows:

1. If GPS reception has never been valid earlier in the flight, the IRS values are used without correction.
2. At present, if GPS reception is lost after a valid complementary-filter correction has been obtained using the procedure described above, but no valid Schuler-oscillation fit has been accumulated as described in [3.] below, the correction factors from the complementary filter at the time that GPS information is lost are reduced by a factor of 0.997 each second, producing an exponential decay back toward the IRS values with decay time constant of about 5.6 min. This time constant can be reset from the “defaults” file.
3. Whenever GPS reception is good, a fit matrix for a least-squares fit to the difference between the GPS and INS groundspeeds is updated, for each component. The errors are assumed to result primarily from a Schuler oscillation, so the three-term fit is of the form $\Delta = c_1 + c_2 \sin(\Omega_{Sch}t) + c_3 \cos(\Omega_{Sch}t)$, where Ω_{Sch} is the angular frequency of the Schuler oscillation (taken to be $2\pi/(5067s)$) and t is the time since the start of the flight. A separate fit is used for each component of the velocity and each component of the position (discussed below under LATC and LONC). The fit matrix used to determine these coefficients is updated each time step but the accumulated fit factors decay exponentially with about 30-min decay constant, so the terms used to determine the fit are exponentially weighted over the period of valid data with a time constant that decays exponentially into the past with a characteristic time of 30 min. This is long enough to determine a significant portion of the Schuler oscillation but short enough to emphasize recent measurements of the correction.

The procedures for accumulating the matrices for the fit are as follows:

Define u_G as the aircraft eastward velocity measured by the GPS and u_I the corresponding velocity measured by the IRS, so that the difference is

$$\delta u = u_G - u_I$$

If Ω_S is the Schuler oscillation period, with $\Omega_S = 2\pi t/T_s$ where $T_s = 5040$ s, τ_u is the time constant for the update (1800 s), t is the time from the start of the flight, and the measurement matrix is A_{ij} , then updated terms of the measurement matrix each sample period ($A'_{i,j}$) are (for δu):

$$A'_{0,1} = A_{0,1} \left(1 - \frac{1}{\tau_u}\right) + \delta u$$

$$A'_{1,1} = A_{1,1} \left(1 - \frac{1}{\tau_u}\right) + \delta u \sin(\Omega_S t)$$

$$A'_{2,1} = A_{2,1} \left(1 - \frac{1}{\tau_u}\right) + \delta u \cos(\Omega_S t)$$

The matrix components $A_{j,0}$ apply to the northward velocity component and so are represented by the same equations with δu replaced by δv . Similar matrices are calculated for latitude θ and longitude ϕ , based on the differences $\delta\theta$ and $\delta\phi$ between GPS and IRS measurements. The information matrix H_{ij} is calculated via

$$H_{i,j} = H_{i,j} \left(1 - \frac{1}{\tau_u}\right) + V_{i,j}$$

where $V_{0,0} = 1$, $V_{0,1} = V_{1,0} = \sin(\Omega_S t)$, $V_{0,2} = V_{2,0} = \cos(\Omega_S t)$, $V_{1,1} = \sin^2(\Omega_S t)$, $V_{1,2} = V_{2,1} = \sin(\Omega_S t) \cos(\Omega_S t)$, and $V_{2,2} = \cos^2(\Omega_S t)$. When the fit is needed, the matrix $H_{i,j}$ is inverted and the result multiplied by the measurement matrix $A_{i,j}$ to get the fit coefficients C_{ij} to use for predicting the results for δu , δv , $\delta\theta$, and $\delta\phi$ via equations like $\delta u = C_{0,1} + C_{1,1} \sin(\Omega_S t) + C_{2,1} \cos(\Omega_S t)$.

4. When GPS data become invalid, if sufficient data (spanning 30 min) have been accumulated, the accumulated fit matrices are inverted to determine the coefficients $\{c_1, c_2, c_3\}$ and then the formula for Δ in the preceding step is used to extrapolate the correction to the IRS measurements while the GPS measurements remain invalid. Doing so immediately could introduce a discontinuity in $\{\text{VNSC}, \text{VEWC}\}$, however, so the correction Δ is introduced smoothly by adjusting $\{\text{VNSC}, \text{VEWC}\}$ as follows: $\Delta' = \eta(\text{VEWC}_0 - \text{GVEW}_0) + (1 - \eta)\Delta$ where Δ' is the sequentially adjusted correction, $(\text{VEWC}_0 - \text{GVEW}_0)$ is the difference preserved from the last time the GPS groundspeed was valid, and $\eta = 0.997 s^{-1}$ is chosen to give a decaying transition with a time constant of about 5.5 min. This has the potential to introduce some artificial variance at this scale and so should be considered in cases where

variance spectra are analyzed in detail, but it has much less influence on such spectra than would introducing a discontinuous transition. Ideally, the current fit and the last filtered discrepancy ($VEWC_0 - GVEW_0$) should be about equal, so this should not introduce a significant change. The current code requires 30 min of valid data to make this fit; otherwise, it assumes that insufficient data have been accumulated for the fit and uses the last determined offset from step 2 if available or the IRS information data if not.

5. When GPS information is again valid, the solution returns to the values of the complementary filter, updating of that filter is resumed, and accumulation of the Schuler-fit coefficients resumes.

Aircraft position

A simpler approach using an exponential filter is possible with aircraft position, as follows:

LAT = latitude measured by the IRS
 LON = longitude measured by the IRS
 GLAT = latitude measured by the GPS
 GLON = longitude measured by the GPS
 VNSC = aircraft ground speed, north component, corrected
 VEWC = aircraft ground speed, east component, corrected

1. Initialize the corrected position at the IRS position at the start of the flight or after any large change ($>5^\circ$) in the IRS position.
2. Integrate forward from that position using the aircraft groundspeed with components $\{VNSC, VEWC\}$. Note that in the absence of GPS information this will introduce long-term errors because it does not account for the Earth's spherical geometry. It provides good short-term accuracy, but the GPS updating in the next step is needed to compensate for the difference between a rectilinear frame and the Earth's spherical coordinate frame and provides a smooth yet accurate track.
3. Use an exponential adjustment to the GPS position, with time constant that is typically about 100 s.^a
4. To handle periods when the GPS becomes invalid, use an approach analogous to that for groundspeed, whereby a Schuler-oscillation fit to the difference between the GPS and IRS measurements is developed and used to extrapolate through periods when the GPS is invalid.

^aspecifically, $LATC += \eta(GLAT - LATC)$ with $\eta = 2\pi/(600s)$

Present Processing Code

See /scr/raf/cjw/builds/bora/raf/nimbus/src/amlib/std/gpsc.c

Reasons For Proposing Changes

The purpose of this memo is primarily to document the procedure, but there are a few minor changes that might be improvements:

1. With present GPS reliability, it may be better to suppress the decay discussed in the box for VNSC/VEWC (#2 in the first box above). This decay allows the merged signal to move back toward the IRS value, but if a good offset has been obtained this is probably not best for short-term drop-outs. I suggest continuing to apply the last offset instead, without decay. This would mean suppressing the section in `gpsc.c` with lines like `dvy[]*=fctrf[]`. The same applies to position, which now is allowed to adjust slowly back to the IRS position; it may be better to retain the last-valid offset.
2. Step 5 above seems to me to be a problem with the current approach because it introduces a discontinuous change from the Schuler-oscillation fit to the complementary-filter fit, without any transition. It would be desirable to use a gradual transition back to the complementary-filtered solution. Specifically, if GPS becomes again good while the Schuler fit is being used, the correction applied to the IRS measurements should be moved toward the complementary-filter solution gradually, with the same exponential filter used to initiate the use of the fit.¹ To move from that previous Schuler-oscillation fit to the complementary-filter fit smoothly, we could use something like `dvy[] = dvy[]*fctrf[]+(1-fctrf[])*filter((gvns-vns), zf[])`. This introduces another time constant (5.6 min) into the filter that otherwise has a 10-min period, but avoids the sharp transition. It compromises the otherwise favorable characteristic of the complementary filter to have no or little effect on the variance spectrum, but `dvy` can be regarded as a correction to `vns` to correct for long-term drift in the IRS, and to the extent that this is true and all variance remains in `vns` then any filtering applied to `dvy` won't have much effect.²
3. (Thanks to Dick Friesen:) There is an error in the filter code. The coefficient b_1 should be $\sqrt{3/2} = 1.227\dots$, not $\sqrt{3}/2=0.866\dots$ as in the code that has been in use since 1992. This doesn't make much difference because the filter with the wrong coefficient still works reasonably although not optimally.

¹Note that there is no need to filter `gvns` and `vns` separately, only their difference, because the filter is linear. We could save two filter steps and arrays: `dvy[]=filter((gvns-vns), zf[])`.

²It may even be acceptable to apply this correction always, instead of switching it on and off when transitions occur, but I haven't investigated this.

Recommendations

1. Use the material in this note as an addition to the Processing Algorithm documentation being developed to replace Bulletin 9 Appendix B.
2. When the GPS signal is lost, suppress drift back to the IRS values; instead, retain the last-determined offset between the IRS and GPS. For example, this would be achieved by omitting the four statements in this code segment from `gpsc.c` starting with line 345 that cause the corrections `{dvy,dvx,dlat,dlon}` to decay with a time constant of about 5.6 min:

```

else
{
/* No good fit accumulated yet, so just let last correction
 * factors slowly decay.
 */
/*      gpsflg = 0.001;
*/
dvy[FeedBack]      *= fctrf[FeedBack];
dvx[FeedBack]      *= fctrf[FeedBack];
dlat[FeedBack]     *= fctrf[FeedBack];
dlon[FeedBack]     *= fctrf[FeedBack];
}
}

```

3. When the GPS signal becomes valid again for GPS measurements, and if a Schuler-fit correction is being used, make a decaying-exponential transition back to the GPS values. (This is already done for position measurements.) This could be implemented by using a different variable `Dvx/Dvy[FeedBack]` to represent the correction, whether determined from the Schuler fit or from the complementary filter. It would then be updated to `dvx/dvy[FeedBack]`, sometimes determined from the complementary filter and sometimes from the Schuler oscillation, depending on validity of the GPS measurements. Code something like the following, to be inserted in place of the two statements that follow “label546:” below line 402 in the code, would accomplish this:

```

// (initialize Dvx, Dvy[FeedBack] to zero // before starting and at line 409)
// then update Dvx to dvx:
Dvx[FeedBack] += (1-fctrf[FeedBack]) * (dvx[FeedBack] - Dvx[FeedBack]);
Dvy[FeedBack] += (1-fctrf[FeedBack]) * (dvy[FeedBack] - Dvy[FeedBack]);
vns[FeedBack] = vns + Dvy[FeedBack];
vew[FeedBack] = vew + Dvx[FeedBack];

```

If this is used, the earlier code that makes a similar transition from GPS-valid to GPS-invalid conditions, located below the commented “`gpsflg = 1.0`” statement at line 331 in `gpsc.c`, should be replaced by the following:

```
dvy[Feedback] = c[0][0] + c[0][1] * sinwt + c[0][2] * coswt;  
(and analogously for the three statements for dvx, dlat, and dlon)
```

4. Although it would not make any change in the output, the coding of the filtering could be simplified by the elimination of two filters because the only filtering that is needed is that applied to the differences (gvns - vns) and (gvew - vew). All four quantities are filtered individually now and the filtered quantities are subtracted to obtain the filtered values of the quantities in parentheses, but the filters are linear so those could be replaced by filters for the two differences to make the routine more elegant.
5. The coefficient b1 in the filter should be corrected to be $\sqrt{1.5}$.

— END —