U.S. Department of Commerce
National Oceanic and Atmospheric Administration
National Weather Service
National Centers for Environmental Prediction
5830 University Research Court
College Park, MD 20740

Office Note 485

doi:10.7289/V5/ON-NCEP-485

COMPARISON OF FINITE DIFFERENCING, TIME SMOOTHING AND SPLINE FITTING ALGORITHMS FOR ESTIMATING AIRSPEED METADATA FROM COARSE-RESOLUTION AIRCRAFT POSITION AND WIND REPORTS

R. James Purser* and Yanqiu Zhu
IM Systems Group, Rockville, Maryland

December 30, 2016

* email: jim.purser@noaa.gov

Abstract

As operational data assimilation procedures become more sophisticated and the need to address and allow for biases and non-Gaussianity in the given data becomes more pressing, we find we need to augment our knowledge about the details concerning the data to include not just the assimilated data values themselves, but also the "metadata" that describe the differing conditions under which the assimilated measurements are made. In the case of aircraft measurements of temperature such relevant metadata used to infer the overall instrument bias might include the altitude, attitude and airspeed of the aircraft. Unfortunately, such parameters are not generally readily available in the reports themselves and must often therefore be inferred indirectly from whatever information *can* be gleaned from the given reports. One good example of this is the airspeed; the attempt to estimate it from successive reported positions together with the measured wind runs into the difficulty that reported positions and times are usually given only to within very crude figures of precision.

The focus of this note is the problem of estimating approximate aircraft motion by applying finite differencing and time-smoothing to the reported approximate positions (it being assumed here that sufficiently precise temporal coordination has been achieved by other means), or fitting these approximate positions, with some slack commensurate with their precision, using an adaptation of the method of "slalom splines". The smoothing method we propose is computationally efficient and scales as a "fast" method with respect to a change of its characteristic time-scale. As we demonstrate, this method usually leads to plausible profiles of aircraft motion which are not obtained by naive finite difference methods and which appear to be as good as, and more robust than, the results given by the more sophisticated (and more costly) spline method.

1. Introduction

The problem we address is one that arises when, for the purposes of better bias correction of aircraft temperature reports (Isaksen et al., 2012; Zhu et al., 2015), we wish to augment the standard profile of measurements with corresponding estimates of the airspeed (which is not one of the standard elements we can obtain from these reports). The position reports are currently given to the nearest hundredth of a degree in longitude and latitude, but some aircraft locations derive from original reports in degrees and minutes, which have less resolution. The time of each aircraft report is given, sometimes to the nearest second, but in many cases only to the nearest minute. Near the ground, where reports are as frequent as every six seconds in order to resolve the greater thermal and dynamical detail seen in atmospheric parameters there, the nominal nearest minute reporting times are therefore often duplicated over and over. Fortunately, the application of the "tension spline" profile-fitting algorithm of Purser et al. (2014) appears to successfully address the problem of separating the reported times into realistic profiles whose temporal data are distinct and monotonic. But the remaining imprecision in the reported horizontal positions results in them exhibiting multiple duplicates, especially close to the ground where the reports are more frequent (implying zero speed!). Even at higher altitudes, with less

frequent reporting, the reported positions remain too imprecise for plausible ground speeds to be extracted from them by simple finite differences alone. The effect is shown for the first part of a real ascending profile in Fig. 1. Another implicit limitation of any interval-averaged piecewise-constant estimates of speed or velocity is that these estimates still fail to provide an unambiguous value at the nodes themselves, where they undergo a discontinuity.

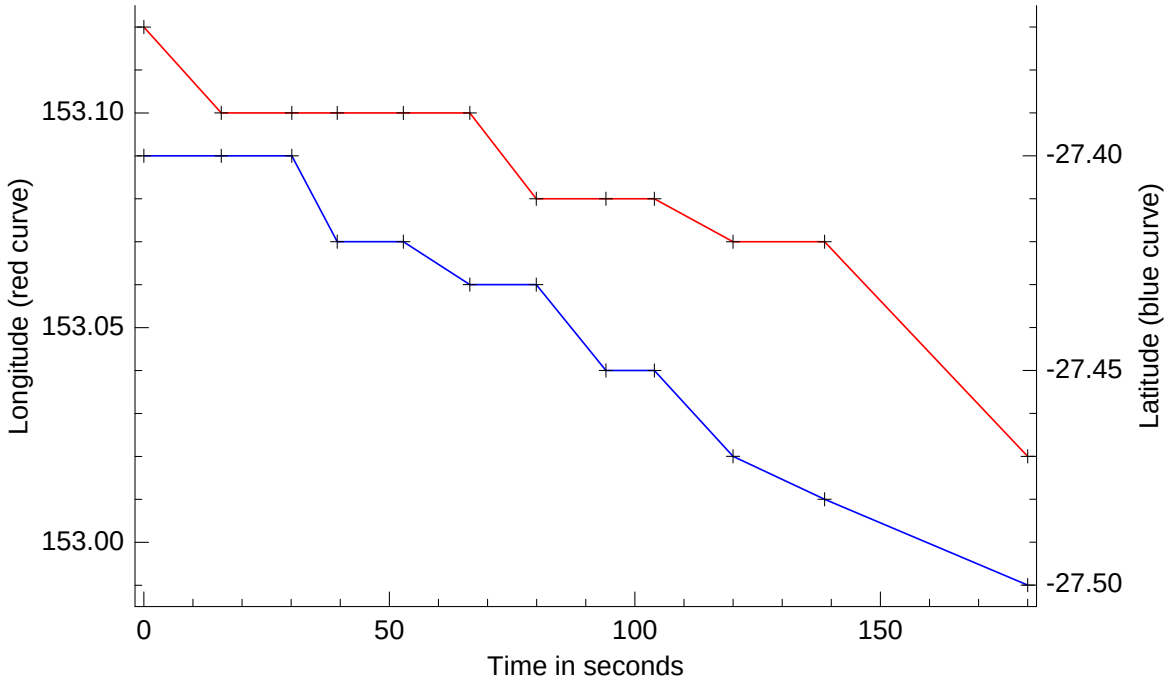

Figure 1. A plot of reported longitudes (red) and latitudes (blue) for the first three minutes of an aircraft ascent, showing the stepped appearance of these position coordinates with time that results from the combination of poor spatial resolution and relatively high frequency of reporting typical of operations close to the ground. A naive application of finite differencing alone to obtain estimates of ground speed will result in very large errors and, in this example, several instances of vanishing derived speed.

A sophisticated variational approach to the reconstruction of aircraft trajectories from incomplete data that employs detailed physical models based on the aerodynamic characteristics typical of reporting aircraft during their ascents and descents has been developed and described by Drüe (2011). For our purposes, we expect that a simpler approach based essentially on geometrical considerations only should be adequate to provide at least a rough estimate of airspeed as one of several potential bias predictor variables that we are investigating. Thus, in order to address the deficiencies of the reported metadata we present here a comparison of relatively simple methods: variants of centered finite differencing; triangular-weighted symmetrical time smoothing of the step-wise velocities or speeds inferred by naive two-point differencing; and a more sophisticated fitting through the finite-width intervals of reported horizontal position coordinates by means of an adaptation of the method of "slalom splines" described in Purser et al. (2014) and Zhu et al. (2015). In each case, the reported wind velocity is subtracted from the inferred aircraft ground-relative motion to obtain estimates of airspeed.

Even after the reported times have been replaced by more reasonable monotonic times via

the aforementioned slalom spline fitting to the reported vertical positions, the replacement of the interval-average two-point finite differencing of Fig. 1 by simple centered differencing does not completely eliminate occurrences of implied zero ground-relative speeds. An alternative procedure of centered finite differencing is tested, where the chosen times before and after the central one are forced to occur where reported positions differ from that of the central time, so that spurious zero velocities cannot so easily occur.

In the case of time smoothing experiments, we continue to first associate with each time interval between reporting nodes of the ascending or descending profile the simplest two-point finite difference estimates of ground-relative velocity as used to produce Fig. 1. An earth-centered cartesian velocity vector avoids any difficulties with the turning of local geographical reference frames with horizontal position. We subtract each reported horizontal wind vector at each node from the ground-relative aircraft vector in each interval to obtain both the air-relative interval-averaged velocity estimate, and hence its corresponding airspeed, within each interval. Both quantities, one a vector and one a scalar, are therefore distributed in time as step-wise constant functions (just as the derivatives of the graphs of Fig. 1 provide step-wise constant estimates of ground-relative motion), and both are subjected to time smoothing in the same way. The time smoothing is performed by a numerically fast algorithm, whereby the smoothed result at each node of the profile becomes the local weighted average of values at nearby times, the temporal weighting function being of symmetric triangular form extending forwards and backwards by some chosen characteristic time, t_c. In the case of the smoothed cartesian air-relative velocity vectors, the corresponding airspeed estimate is taken, at the nodes, as the final computational step. In general, it differs slightly from the smoothed estimate of airspeed obtained by smoothing the piecewise-constant airspeed function directly.

Finally, we test an adaptation of the slalom spline method that finds the cubic spline of least "energy of bending" that can be threaded through all of the finite-width precision-tolerance "slalom gates" of reported longitude and the similarly formulated spline threading the corresponding precision-tolerance slalom gates of reported latitude. Using the same profile data of Fig. 1, but replacing those longitude and latitude point values with finite-width acceptance gates of width 0.01 degrees in both longitude and latitude the minimum energy slalom cubic splines that thread these gates (and all the other ones of the profile not shown) we obtain the graphs of position shown in Fig. 2.

The next section will describe these formulations in greater detail and section 3 will report on results. The implications for operational bias correction are discussed in the concluding section.

2. Estimates of airspeed by the different methods

We assume the profile occupies a period $[t_0, t_1]$ of time, t, within which there are several other nodes where imprecise position data are provided. If λ is longitude, ϕ the latitude available at each data node of the profile, it is easy to compute the three-dimensional cartesian position vector (in earth-radius units at this stage),

$$\boldsymbol{z} = [\cos(\phi)\cos(\lambda), \cos(\phi)\sin(\lambda), \sin(\phi)] \,, \tag{2.1}$$

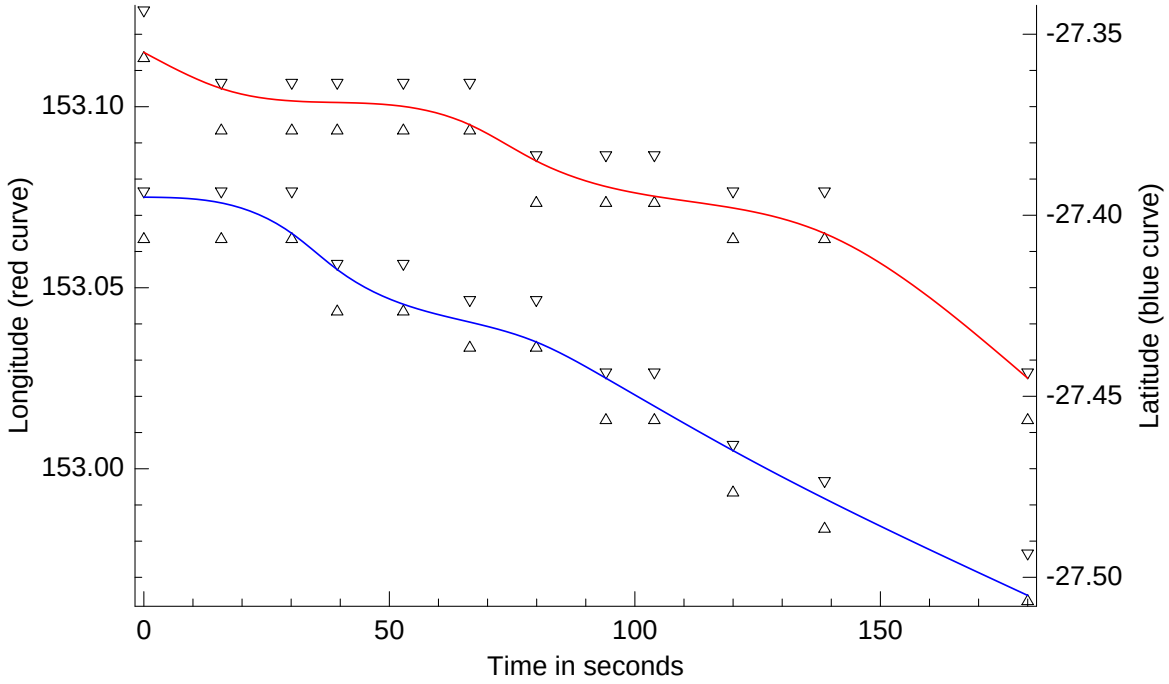

Figure 2. Like Fig. 1 except that each longitude and latitude report is fitted, within a tolerance modeled by a "slalom gate" of width 0.01 degrees in both longitude and latitude, by the cubic spline that has the minimum bending "energy" of all such splines that thread the given gates.

and hence an estimate of the corresponding three-vector of cartesian velocity, $\boldsymbol{v}$ within each interval by simple two-point finite differencing from the consecutive $\boldsymbol{z}$:

$$\boldsymbol{v}_{pq} = R\frac{\delta \boldsymbol{z}_{pq}}{\delta t_{pq}}, \tag{2.2}$$

where

$$\delta \boldsymbol{z}_{pq} \equiv \boldsymbol{z}_q - \boldsymbol{z}_p \tag{2.3a}$$
$$\delta t_{pq} \equiv t_q - t_p, \tag{2.3b}$$

for nodes labeled p and q and R is the radius of Earth. The motivation for adopting cartesians in the description of horizontal positions and velocities was to avoid any complicating influence of the curvature of the geographical coordinates in these experiments, but this choice turned out to be not strictly necessary for the actual data used.

(*a*) *Estimating airspeeds by centered finite differencing*

Except at the ends of a profile where we have no other choice, we can improve these velocity estimates (for smooth trajectories, at least) at an interior node by the weighted mean of the straddling interval estimates taken from the consecutive nodes, p, q and r:

$$\overline{\boldsymbol{v}}_q \approx \frac{\boldsymbol{v}_{pq}\delta t_{qr} + \boldsymbol{v}_{qr}\delta t_{pq}}{\delta t_{pr}}. \tag{2.4}$$

The cartesian wind vector,

$$\boldsymbol{u} = \boldsymbol{x}u + \boldsymbol{y}v, \tag{2.5}$$

where u and v are the conventional reported wind vector components and

$$\boldsymbol{x} = [-\sin(\lambda), -\cos(\lambda), 0]\,, \tag{2.6a}$$
$$\boldsymbol{y} = [-\sin(\phi)\cos(\lambda), -\sin(\phi)\sin(\lambda), \cos(\phi)]\,, \tag{2.6b}$$

is subtracted from these ground-relative motion estimates to arrive at an estimate of air-relative velocity, and hence airspeed. The following graphs of these naive estimated airspeeds are always shown in magenta. however, as we have noted, the reported position data near the ground often contain duplicated consecutive positions so, to avoid obtaining spurious zero-velocity estimates, we also consider the centered finite difference method in which the two straddling time intervals are no longer restricted to involving consecutive nodes only, but are extended just far enough to more distant nodes (whenever possible) to ensure that the straddling reported positions differ from the central position in the application of (2.4). Throughout most of each profile the resulting calculations do not differ from the simple centered differencing, but where they do, we show the resulting portion of each graph of the corresponding airspeed estimate in red.

(*b*) *Estimating airspeeds by time smoothing basic interval-mean estimates*

For the reasons discussed above, the profile of $\boldsymbol{v}$ resulting from the finite differencing of the inadequately resolved positions leads to significant and obvious errors (as the graphs will show). But the errors in the basic interval-mean $\boldsymbol{v}$ profile will at least tend to compensate locally, and should therefore respond well to a remedial time-smoothing of them. We could choose to run a symmetric moving-average (uniform weighting, defined width) smoother through the data, which would generally improve the results and would lead to a continuous, rather than stepped, profile. But it is also not too difficult to go further and use a triangular weighting function, which should be even better at calming the wilder excursions of the $\boldsymbol{v}$ estimates. The generic unnormalized triangular weighting function can be defined:

$$W(t,t') = \begin{cases} t_c - |t'-t|, & : \quad |t'-t| < t_c \\ 0, & : \quad \text{otherwise} \end{cases}\,, \tag{2.7}$$

where t_c is the characteristic time scale, or half-width, of the smoother. For a given target time, t, define functions,

$$t_-(t) = \max(t - t_c, t_0) \tag{2.8a}$$
$$t_+(t) = \min(t + t_c, t_1), \tag{2.8b}$$

which describe the limits of the weighted average that affects target time, t. Then we can smooth the profile $\boldsymbol{v}(t)$ in $[t_0, t_1]$ to get

$$\overline{\boldsymbol{v}}(t) = \frac{\int_{t_-(t)}^{t_+(t)} W(t,t')\boldsymbol{v}(t')\,dt'}{\int_{t_-(t)}^{t_+(t)} W(t,t')\,dt'}. \tag{2.9}$$

Note that the smoothing weight, W, becomes effectively clipped for any target t closer than t_c to either end of the profile period. We provide the algebraic details of the computationally efficient algorithm that accomplishes a smoothing of this kind in the appendix.

In practice, it is generally the relative motion of the aircraft with respect to the surrounding air that we are interested in smoothing. For this, we just adjust the interval-averages $\boldsymbol{v}_{pq}$ by subtracting from each the average reported wind vector from each end of the interval. Thus, as we have discussed earlier, we can replace the aircraft earth-relative motion vector $\boldsymbol{v}_{ab}$ associated with each interval $[t_a, t_b]$ used in our algorithm to construct the integrals $\boldsymbol{J}$ and $\boldsymbol{K}$ by the corresponding air-relative vector,

$$\boldsymbol{v}'_{ab} = \boldsymbol{v}_{ab} - \boldsymbol{u}_{ab}, \tag{2.10}$$

where $\boldsymbol{u}_{ab}$ is the average wind vector in earth-centered cartesians, computed for that interval at the aircraft's location from reported winds at the relevant nodes a and b. The rest of the analysis proceeds as before. The smoothed aircraft motion vectors relative to the air can be converted back to conventional geographical reference components, eastwards and northwards U' and V', by:

$$U' = \boldsymbol{x}\cdot\overline{\boldsymbol{v}'} \tag{2.11a}$$

$$V' = \boldsymbol{y}\cdot\overline{\boldsymbol{v}'}. \tag{2.11b}$$

The graphs of airspeed, $\sqrt{U'^2 + V'^2}$ derived from smoothing the velocity vectors will be shown in blue; the graphs of the directly smoothed piecewise-constant airspeed estimates, $|v'_{ab}|$, will be shown in green.

(*c*) *Estimates using splines*

The final method of estimating the airspeed from profiles of the imprecise position and wind reports is to use cubic spline functions of time to fit, separately, the longitude and latitude, each within a tolerance "gate" centered on the reported coordinate but having a finite acceptance width guided by the reported precision of these position reports. In Purser et al. (2014) we constructed splines with a similar finite tolerance property except, in that case, the slack was purely in the direction of the independent variable, time, whereas in the present adaptation of that so-called "slalom spline", the time remains precise and the slack resides in the dependent variable (longitude or latitude in this case). Nevertheless, the same general principle applies: thread the spline through the "gates" defined by a set of pairs of opposing inequality constraints (defining a family of "feasible" splines) and, amongst these feasible splines choose the one that minimizes the integral of the square of the spline's second-derivative (an analogue of the spline's "elastic energy"). In Purser et al. (2014) and Zhu et al. (2015), the family of splines deemed appropriate to the problem we were dealing with was of the "tension spline" variety (Schweikert 1966), but in the present case, there is no need to require splines more sophisticated than the ordinary cubic variety, which are also slightly simpler to compute.

The spline-fitted aircraft trajectories are continuous longitude- and latitude-valued functions of time, from which it is straightforward to recover the ground-relative velocities at any points of the profile. At the reporting times, we can again subtract the measured wind velocities to recover estimates of airspeed, but for graphing purposes, we can obtain continuous estimates of

airspeed by subtracting spline-fitted wind velocities, and thence the implied airspeed estimates as a smooth graph. In all the graphs we show these spline estimates in brown.

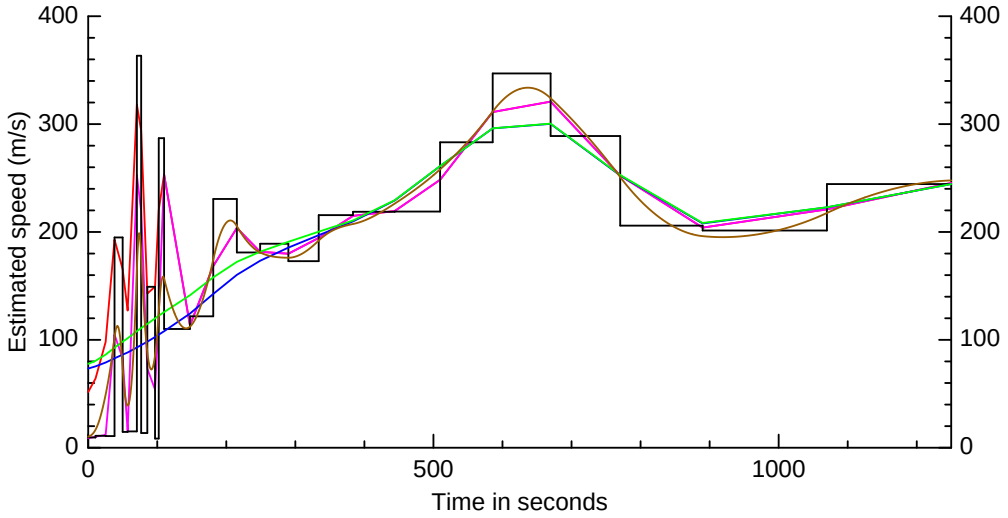

Figure 3. A plot of the full set of estimates of aircraft airspeed derived from the sequence of the coarsely resolved (hundredth of a degree) geographical coordinates at each node of a typical ascending profile. The speeds derived using simple two-point differencing of reported aircraft locations form the piecewise constant graph resembling a histogram, shown in black. The magenta curve shows the airspeed graph corresponding to ordinary centered differencing estimate of velocity defined in (2.4). At the left, corresponding to the early phase of the ascent close to the ground, the reports are very frequent and the speeds are relatively small. The alternative centered differencing estimates which force the use of distinct positions are shown in the red graph that splits from the magenta there. A less cluttered view of the graphs for just the first 250 seconds of this ascent is shown in Fig. 4. The triangular-weighted time-smoothing of the raw piecewise-constant velocity vector function with a time constant of 180 seconds, followed by conversion of the resulting smoothed vectors to speeds, produce the blue curve. If we apply the same smoothing method to the speeds directly, we obtain the almost identical green curve. Finally, the brown curve shows the slalom spline estimate of airspeed when the "gates" threaded by the splines in longitude and latitude separately are of width 0.01 degrees. The details of these plots in the cluttered region near the ground, corresponding to the left side of the figure, are better seen in the expanded view of this region shown in Fig. 4.

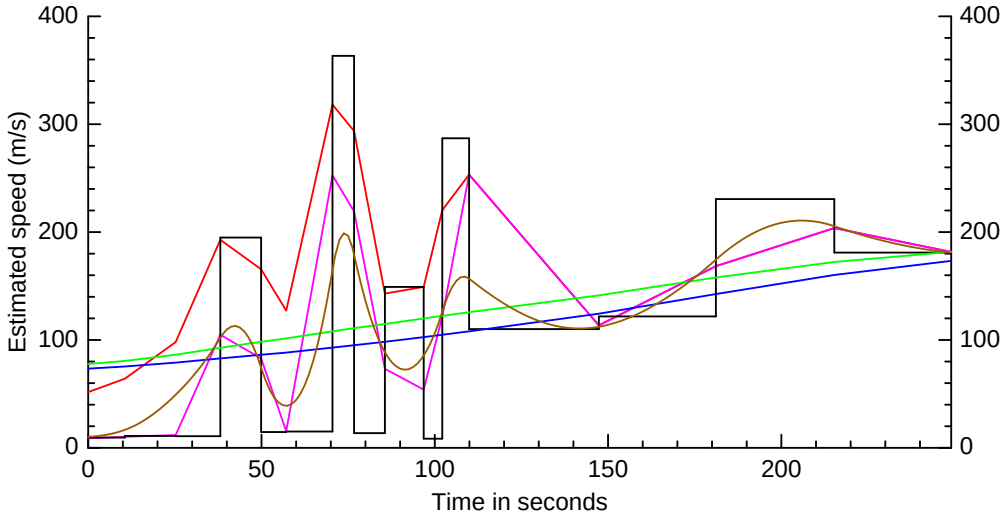

Figure 4. An expanded view of just the first 250 seconds of the same ascent shown in Fig. 3 plotting the same graphs as shown in Fig. 3. Compared to ordinary centered finite difference estimates (magenta), the red graph shows a slightly less wildly fluctuating estimate obtained when the straddling times used are required to have reported positions that differ from the position of the central target at each node of the profile. However, the improvement is small and inadequate. In this expanded view we more clearly see the slight deviation that occurs between the graphs of airspeed implied by time-smoothed velocity (blue) and the more directly obtained estimates of time-smoothed airspeed itself (green). Finally, the large excursions of the spline-based estimate (brown) are clearly seen in this part of the profile.

3. Numerical experiments and results

(*a*) *Single-aircraft experiments*

A typical ascending profile is shown in Fig. 3 with graphs of all the methods plotted for comparison. The black stepped graph is just the plot of interval-averaged airspeeds $|\boldsymbol{v}'_{ab}|$ derived from two-point finite differences of reported position combined with the interval-mean reported wind veclocity correction. The vertical sides of this histogram-like graph indicate the locations of all the nodes of the profile. Clearly, these naive estimates cannot directly supply unambiguous estimates of airspeed at the nodes themselves, but a trivially more complicated centered modification of them, using (2.4), does provide such estimates at the nodes. These estimates are shown as the magenta graph. At the far left of the graph, we have more frequent reporting and the finite difference estimates clearly become quite useless, owing to the poor precision of the reported longitudes and latitudes from which they are derived. Even the modification to the centered differencing that forces the selected locations at either side of the target time to be distinct from the target's reported location does not substantially improve the quality of the resulting airspeed estimates, as we see in this figure and the expanded view of the same graphs at the first approximately 250 seconds shown in Fig. 4. In fact, in *all* of the profiles we have tested, these finite difference estimates show unrealistic large-amplitude excursions, so we shall not consider these failed methods further.

The two graphs employing time smoothing merge into a single green curve for most of the ascent of Fig. 3 but, at the very early times, there is a slight separation between the results of smoothing velocity (blue) and smoothing speed directly (green), as we see more clearly in the

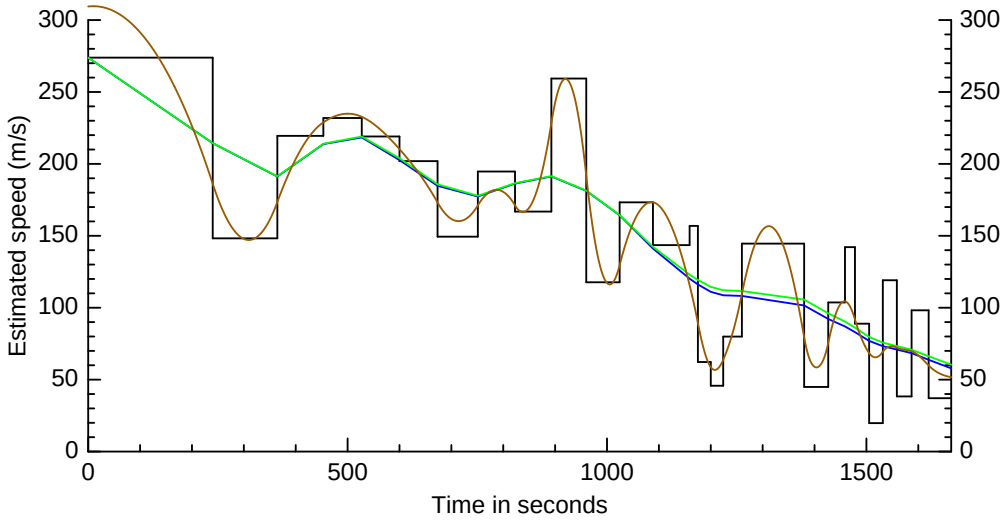

Figure 5. Like Fig. 3 except for a descending profile and without the centered finite differences shown. The unrealistic excursions of the spline (brown), and to a lesser degree, the time-smoothed plots (green), suggest that the parameters should be adjusted to provide a greater degree of smoothing than is used here.

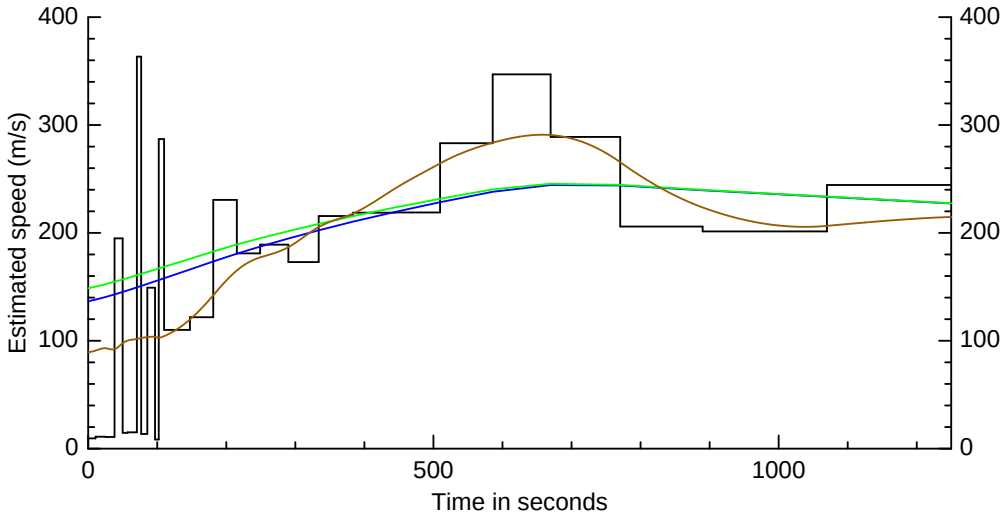

Figure 6. Like Fig. 3 but with a larger smoothing parameter, $t_c = 600s$ for the time smoothing (green and blue curves), and a broader slalom gate width of 0.04 degrees for the spline (brown).

expanded version, Fig 4. Both these curves also remain close to the estimate of airspeed derived from the application of the slalom spline method, here applied with a gate width (in longitude and latitude) of 0.01 degrees, except near the ground where the spline estimates fluctuate with especially unrealistically large amplitude (see Fig. 4 again). But even in the middle portion of the ascent where all the estimates are roughly consistent, we are forced to question their correctness since the magnitudes of these airspeed estimates are roughly sonic (the speed of sound being around 300 ms^{-1} here) and some other airspeeds are too slow to be consistent with

safe flight.

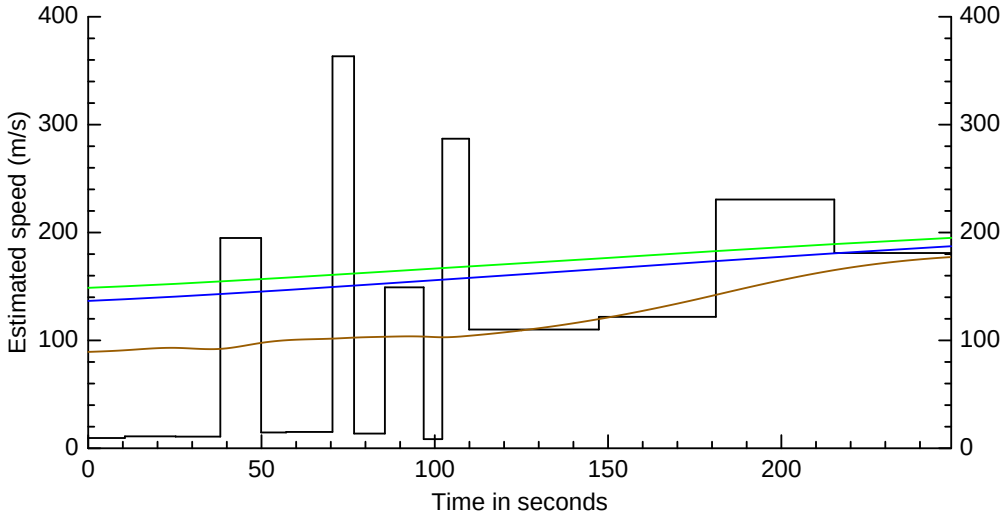

Figure 7. Like Fig. 6 but showing detail at the first 250 seconds.

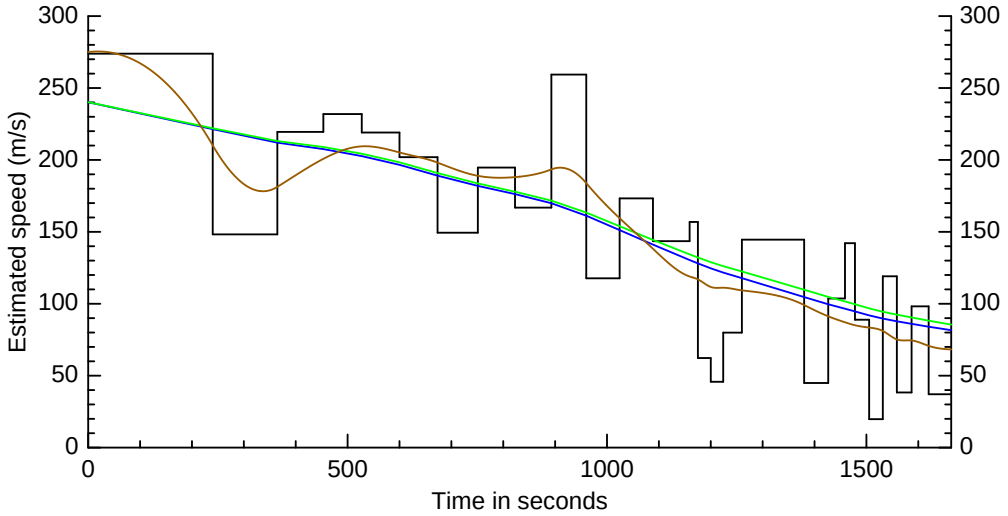

Figure 8. Like Fig. 5 but with time smoothing parameter $t_c = 600s$ (green and blue curves) and gate width 0.04 degrees for the slalom spline (brown).

Another profile, this time descending and not cluttered by the centered differencing results, is shown in Fig. 5 for the same characteristic smoothing time $t_c = 180$ s (green and blue curves) and slalom gate width of 0.01 degrees (brown curve) for the spline solution. The exuberant excursions of the spline solution adhere much too closely to the raw histogram-like interval-averaged speed graphs in these figures, suggesting that the provided reports and inferred times still contain errors larger than are accounted for by simply setting the gate width to the nominal

precision of the position reports.

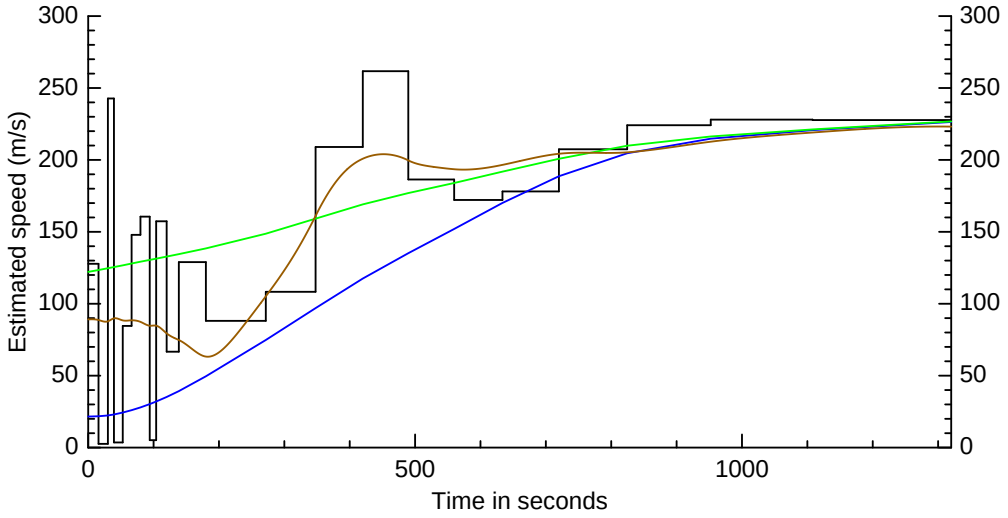

Figure 9. Like Fig. 6 with $t_c = 600$ s and gate width 0.04 degrees, but for a strongly curved ascending profile, showing how, in such a case, the smoothing of interval-averaged velocities of the profile can lead to a severe under-estimation of the deduced airspeeds. A similar speed deficit is seen using the slalom spline which fits the slack position coordinates with an unrealistically tight curvature.

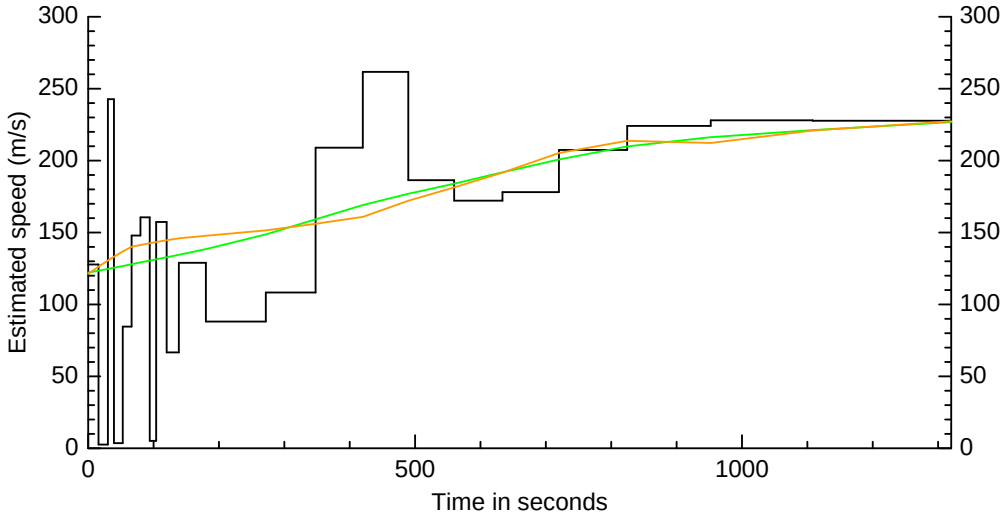

Figure 10. The same example as in Fig. 9 but comparing only the triangular-weighted time smoothing of interval-averaged estimates of airspeed with the nearest equivalent smoothing ($t_c = 424.264$ s) that employs a uniform weighting (orange curve), which always results in a noisier graph.

While it is not practical to address the possible imprecision of the assumed times, we can, in compensation, artificially and arbitrarily broaden the width of the assumed slalom gates. The same excessive waviness, though less severe, appears in the time-smoothed graphs, again suggesting that the responsible parameter, in this case t_c, be increased. Therefore we repeat the experiments, but this time with $t_c = 600$ s (ten minutes) and slalom gate width 0.04 degrees.

For the ascending example, the new plots are shown, in the same formats as Figs. 3 and 4 respectively, in Figs. 6 and 7. The new plot for the same descending profile of Fig. 5 is shown in Fig. 8. In these new plots, the profiles using the time smoothing of velocity (blue) and speed (green) are much improved and, in these plots at least, mutually consistent. The plots of the estimates of airspeed given by the splines remain unsatisfactory, however, with persisting high and low excursions that do not seem realistic.

Although in the two cases we have looked at so far, the time smoothing of velocity or of the speed directly seem to give very similar results for the respective airspeed estimates, this is not always so. In Fig 9 we show an example of an ascent in which the airspeeds estimated from a smoothing of velocities (blue) give values very much lower (and less realistic) than the outcome of time smoothing the interval-average airspeed estimate directly (green). In this case, the aircraft executes a fairly tight turn soon after take-off so that the velocity vectors turn on a time scale commensurate with the smoothing time scale, $t_c = 600$ s. Smoothing the vectors is clearly inappropriate in such cases. It is also notable that the slalom spline applied with the newly enlarged gates, width 0.04 degrees, also gives too much slack to the effective locations at each reporting time, resulting in a "corner-cutting" trajectory estimate whose implied speed is, again, far too small at around 200 s into the ascent. It is on the basis of this and similar profile results that we are inclined to suggest that the best practical method for recovering plausible airspeeds from the coarse position data we are provided with is to use a time smoothing of the airspeeds with this larger choice of time constant, $t_c = 600$ s.

One final single-aircraft experiment will compare the triangular-weighted time smoothing of airspeed of the previous example, with the same time constant, $t_c = 600$ s, with the uniformly-weighted ("centered moving average") smoother of (A.8), but with the half-width of the support of the weighting function reduced by a factor of $\sqrt{2}$, i.e., $t_c = 424.264$ s for the moving average, in order to maintain the same normalized second moment (or same "radius of gyration") as in the triangular-weighted case. The two curves are shown in Fig. 10 with the uniform-weighted case plotted in orange. While both plots are broadly similar, the uniform-weighted smoother is distinctly less smooth (this is true of all the other profiles we have looked at).

(*b*) *Statistical results from different reporting types*

As yet we have no routine access to the "true" airspeeds for aircraft data we are using; experiments using some reports where three decimal place latitude and longitude positions were provided allowed us to artificially degrade these data and show that the time smoothing method was at least giving consistent answers, but the time precision (the effective "weak link" for getting speed estimates) of these data remained inadequate for a validation using direct estimates of the true airspeed by finite differencing alone. Some of the MDCRS data now becoming available (Bradley Ballish, personal communication) combine the more precise positional data with reliable nearest-second time reporting, and we plan to validate our proposed time smoothing method with samples of these data in a future study. Meanwhile, we can at least get some idea about the plausibility of the deduced airspeed estimates from our time-smoothing strategy by examining the histograms of data counts against deduced airspeed, and seeing the proportion of these counts that lie clearly outside the typical operating parameters of reporting aircraft. Figures 11 and 12 plot the data counts corresponding to these raw deduced airspeeds (based on the triangular weighting with a time parameters of 600 s). If we take the

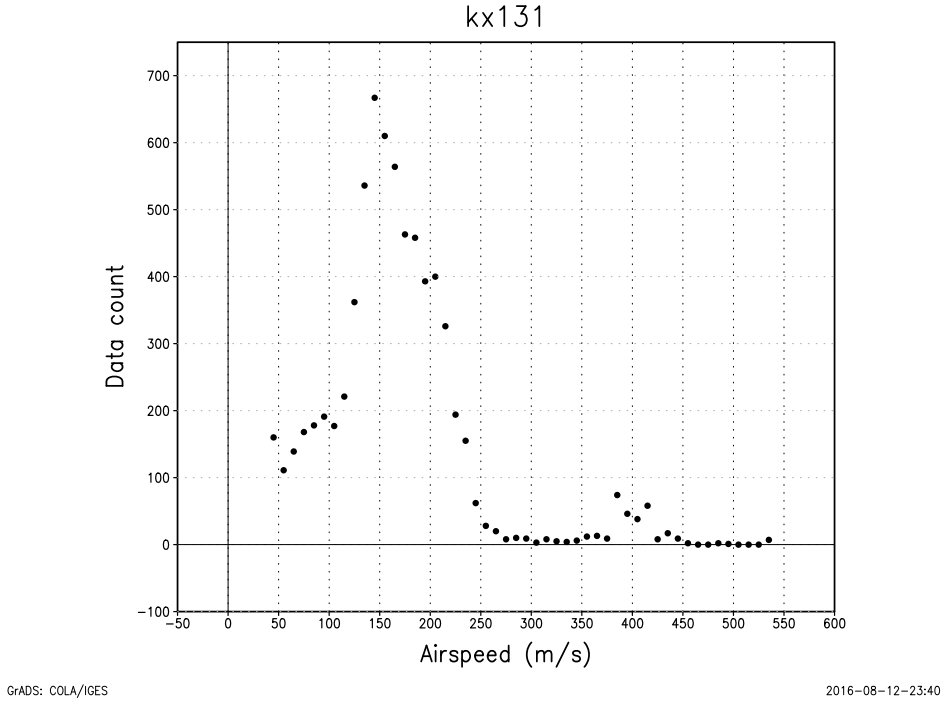

Figure 11. Histogram of counts of data obtained from a period of one month from the aircraft of type 131 (Non-US AMDAR sensible temperature) plotted against raw (uncapped) deduced airspeed in 10 ms^{-1} intervals from the triangular-weighted 600-second time-smoothing of speeds.

lowest plausible airspeed to be 70 ms^{-1} and the highest to be 300 ms^{-1}, we see that there are indeed a few retrieved airspeed estimates that violate these limits, but they are sufficiently few to suggest that, provided we cap the retrieved values to enforce these limits, the results we obtain might still be adequate as inputs to any bias-estimation procedures we may wish to implement (given that we have no other better airspeed estimates to work from).

For the data type 131 (non-US AMDAR) we find approximately 5.9% of the data fall below the 70 ms^{-1} limit; 4.6% above 300 ms^{-1}. For the data type 133 (MDCRS) , the percentages of violations are: 2.7% below 70 ms^{-1} and 2.7% above 300 ms^{-1}. While these percentages are clearly higher than we would wish, they are probably sufficiently small to suggest that our crude but efficient method for extracting something approximating the airspeed is at least of some value in the context for which it is intended to be used (i.e., to provide a physically meaningful bias predictor variable). We expect that future enhancements of the aircraft data quality control will lead to further improvements in the airspeed estimates that our method provides.

4. Discussion and Conclusion

Faced with the practical problem of deducing plausible airspeeds from wind and coarsely-resolved position and time reports provided by aircraft during their ascents and descents, we have compared several simple and numerically efficient estimation procedures. The methods

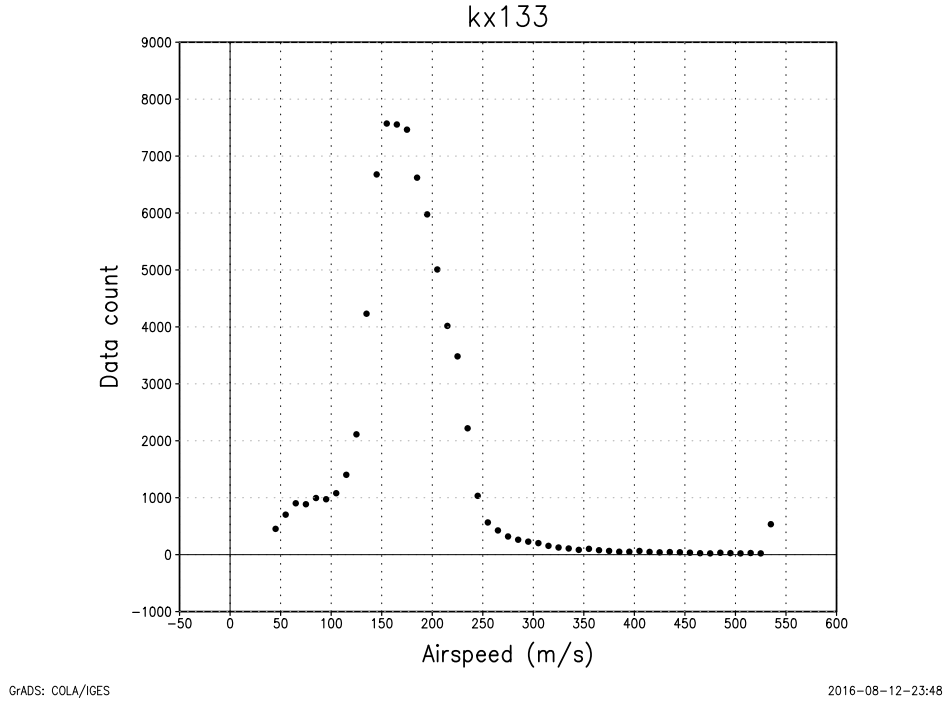

Figure 12. Like Fig. 11, but from the aircraft of type 133 (MDCRS sensible temperature).

based on centered finite differences of the reported positions, even when provisions were made to avoid using duplicated positions in the differencing, were immediately found to be intolerably noisy. The other practical options include time-smoothing the piecewise-constant interval-average estimates of velocity or airspeed, and fitting the positional data, with a built-in slack reflecting the imprecision of these reports using cubic slalom splines. We have described in detail a fast but simple triangular-weighted time smoothing algorithm which appears to accomplish our objective when the time constant t_c measuring the half-width of the triangular weight profile is set to around ten minutes, and is applied preferably to the speeds rather than the velocities. We find this method to be slightly superior to the application of an ordinary symmetric moving average filter (which is also fast and simple). We have compared these methods to an adaptation of the slalom splines developed by Purser et al. (2014) to treat a different (vertical) aspect of the aircraft trajectory-estimation procedure, but, in the present case, we find (somewhat surprisingly) that, owing to what is presumed to be a combination of residual errors in the reported positions and the presumed reporting times, we do not retrieve airspeed estimates as plausible as those we attain with the simpler time-smoothing method.

The adopted time-smoothing technique with its somewhat arbitrary choice of a constant ten-minute smoothing parameter, t_c, can easily be made more general when this parameter is allowed to depend upon the phase of flight and the altitude; this would not jeopardize the formal assignation of the algorithm as being a "fast" one. Beyond such minor adaptations of the proposed method, which essentially treats the problem as a purely goemetrical one, a further improvement in the accuracy of the estimate of airspeed, especially at lower altitudes of the

asents and descents, would most likely require the incorporation of additional information about typical flight characteristics, which is the approach adopted and dealt with in a very thorough manner by Drüe (2011). One further development in this direction is to take, as the input for the time-smoothing, the interval estimate of the inferred airspeed minus some altitude- and aircraft-dependent "standard profile" (and, of course, to then add the smoothed result back onto this profile). This would alleviate the defect visible in some of our plots near take-off and landing, where the airspeed is being over-estimated in the smoothing process, which mixes in the higher speeds occurring at the slightly higher altitudes. The standard airspeed profiles need not be very sophisticated, but must be smooth functions of altitude, to provide this benefit. We plan to explore this approach in future.

The method by which we have achieved a formally "fast" algorithm for smoothing with a kernel of triangular shape can be extended to other kernels that belong to the family of what are known as "Cardinal B-splines" (de Boor 1978). These constitute a family of piecewise-polynomial bell-shaped functions whose higher-order members closely resemble Gaussian functions (Schoenberg, 1946; 1973). For this reason, fast convolution smoothing algorithms (but in space instead of time) based on these functions might be of interest as alternatives to the recursive filters (Wu et al. 2002, Purser et al. 2003) used in the generation of practical covariance operators for efficient implementations of variational data assimilation. This is especially relevant now since the attribute of possessing a compact support which these B-spline functions enjoy, but is not shared by the recursive filters used in present operations, is an attractive asset for contemporary massively parallel computer architectures. This will be elaborated in a future report.

Acknowledgments

This note benefitted from stimulating discussions with Drs. Bradley A. Ballish and Andrew Collard. We are grateful to Dr. Chris Hill for providing sample aircraft reports with enhanced positional precision and to Bradley Ballish for his insightful reviews of this note, and for his efforts to acquire MDCRS data for future studies with both temporal and positional high precision.

Appendix A

Efficient smoothing with triangular weighting

In practice, it is convenient to break each integral of (2.9) into two smooth portions: the part "−" with $t' < t$ and the part "+" with $t' > t$. To this end we define new weights:

$$W^-(t,t') = \begin{cases} t_c - t + t', & : \quad t - t_c < t' \le t \\ 0, & : \quad \text{otherwise} \end{cases} , \tag{A.1}$$

and

$$W^+(t,t') = \begin{cases} t_c + t - t', & : \quad t \le t' < t + t_c \\ 0, & : \quad \text{otherwise} \end{cases} . \tag{A.2}$$

Now, (2.9) becomes equivalent to:

$$\overline{\boldsymbol{v}}(t) = \frac{\int_{t_-(t)}^{t} W^-(t,t')\boldsymbol{v}(t')\,dt' + \int_{t}^{t_+(t)} W^+(t,t')\boldsymbol{v}(t')\,dt'}{\int_{t_-(t)}^{t} W^-(t,t')\,dt' + \int_{t}^{t_+(t)} W^+(t,t')\,dt'}. \tag{A.3}$$

Define the (vector-valued) integrals:

$$\boldsymbol{J}(t) = \int_0^t \boldsymbol{v}(t')\,dt' \tag{A.4a}$$

$$\boldsymbol{K}(t) = \int_0^t \boldsymbol{v}(t')t'\,dt'. \tag{A.4b}$$

Now the advantage of splitting (2.9) into the smooth integrals (A.3) emerges as we observe that:

$$\int_{t_-}^{t} W^-(t,t')\boldsymbol{v}(t')\,dt' = (t_c - t)(\boldsymbol{J}(t) - \boldsymbol{J}(t_-)) + \boldsymbol{K}(t) - \boldsymbol{K}(t_-) \tag{A.5a}$$

$$\int_{t}^{t_+} W^+(t,t')\boldsymbol{v}(t')\,dt' = (t_c + t)(\boldsymbol{J}(t_+) - \boldsymbol{J}(t)) + \boldsymbol{K}(t) - \boldsymbol{K}(t_+) \tag{A.5b}$$

$$\int_{t_-}^{t} W^-(t,t')\,dt' = (t_c - t)(t - t_-) + \frac{(t^2 - t_-^2)}{2} \tag{A.5c}$$

$$\int_{t}^{t_+} W^+(t,t')\,dt' = (t_c + t)(t_+ - t) + \frac{(t^2 - t_+^2)}{2}, \tag{A.5d}$$

(where t_- and t_+ are implicitly taken to be the values of these functions at this particular chosen argument t). After we have gathered up the terms, and made the convenient substitutions,

$$\boldsymbol{J}_t \equiv \boldsymbol{J}(t) \tag{A.6a}$$

$$\boldsymbol{J}_\pm \equiv \boldsymbol{J}(t_\pm) \tag{A.6b}$$

$$\boldsymbol{K}_t \equiv \boldsymbol{K}(t) \tag{A.6c}$$

$$\boldsymbol{K}_\pm \equiv \boldsymbol{K}(t_\pm), \tag{A.6d}$$

we can express the result for smoothed value, $\overline{\boldsymbol{v}}(t)$, as:

$$\overline{\boldsymbol{v}}(t) = \frac{t_c\left[\boldsymbol{J}_+ - \boldsymbol{J}_-\right] + t\left[\boldsymbol{J}_+ - 2\boldsymbol{J}_t + \boldsymbol{J}_-\right] - \left[\boldsymbol{K}_+ - 2\boldsymbol{K}_t + \boldsymbol{K}_-\right]}{t_c\left[t_+ - t_-\right] + t\left[t_+ - 2t + t_-\right] - \left[t_+^2 - 2t^2 + t_-^2\right]/2}. \tag{A.7}$$

For comparison, the piecewise-constant, or "centered moving-average" smoothing of $\boldsymbol{v}$ whose weighting function has the same support, $[t_-, t_+]$, is given using the same notation simply as:

$$\overline{\boldsymbol{v}}(t) = \frac{\left[\boldsymbol{J}_+ - \boldsymbol{J}_-\right]}{\left[t_+ - t_-\right]}. \tag{A.8}$$

For the case where we assume $\boldsymbol{v}$ is piecewise constant in the given intervals of the profile, $\boldsymbol{J}(t)$ becomes a piecewise-linear and $\boldsymbol{K}(t)$ a piecewise-quadratic function in those intervals. Both functions are easily computed and stored at the nodes of the profile, allowing $\boldsymbol{J}$ to be

interpolated *exactly* to any target interior to an interval $[t_a, t_b]$ by linear interpolation from the corresponding node values, $\boldsymbol{J}_a$ and $\boldsymbol{J}_b$. But since $\boldsymbol{v}$ is constant, say $\boldsymbol{v}_{ab}$, within the interval, (A.4b) implies that the linear interpolation estimate of the interior target value of $\boldsymbol{K}(t)$ has an error, but one of particularly simple quadratic form:

$$\boldsymbol{K}(t) = \boldsymbol{K}_a + \frac{\boldsymbol{K}_b - \boldsymbol{K}_a}{t_b - t_a}(t - t_a) - \frac{(t - t_a)(t_b - t)\boldsymbol{v}_{ab}}{2}, \tag{A.9}$$

which makes it almost equally easy to also interpolate $\boldsymbol{K}(t)$ exactly. Hence, we are able easily and efficiently to compute the exact time-smoothed solutions, $\overline{\boldsymbol{v}}(t)$, for each node of the profile.

The same algorithm can be applied to smooth a distribution of piecewise-constant scalars, such as inferred airspeeds, $|\boldsymbol{v}'|$, without any change beyond the replacement of all the vector quantities by their corresponding scalar magnitudes. As we shall see, the smoothing of scalar aircraft speeds relative to the air, has some decided advantages in cases where the aircraft track is strongly curved.

References

de Boor, C. 1978 *A Practical Guide to Splines.* Springer, New York. 392 pp.

Drüe, C., 2011 Reconstruction of aircraft trajectories from AMDAR weather reports. *J. Atmos. and Ocean Tech.,* **28,** 921–932.

Isaksen, L., D. Vasiljevic, D. Dee, and S. Healy 2012 Bias correction of aircraft data implemented in November 2011. *ECMWF Newsletter,* **131** 6–6.

Purser, R. J., Y. Zhu, and B. A. Ballish 2014 Recovery of Aircraft Vertical Motion Profiles from Incomplete Data - an Application of the Method of Splines. NOAA/NCEP Office Note 480.

Purser, R. J., W.-S. Wu, D. F. Parrish, and N. M. Roberts 2003 Numerical aspects of the application of recursive filters to variational statistical analysis. Part I: Spatially homogeneous and isotropic Gaussian covariances. *Mon. Wea. Rev.,* **131,** 1524–1535.

Schoenberg, I. J. 1946 Contributions to the problem of approximation of equidistant data by analytic functions. *Quart. Appl. Math.,* **4,** 45–99, 112–141.

Schoenberg, I. J. 1973 *Cardinal Spline Interpolation. CBMS-NSF Regional Conference Series in Applied Mathematics.* SIAM, pp. 119

Schweikert, D. G. 1966 An interpolation curve using a spline in tension, *J. Math. Phys.,* **45,** 312–317.

Wu, W.-S., R. J. Purser, and D. F. Parrish 2002 Three-dimensional variational analysis with spatially inhomogeneous covariances. *Mon. Wea. Rev.,* **130,** 2905–2916.

Zhu, Y., J. C. Derber, R. J. Purser, B. A. Ballish, and J. Whiting 2015 Variational correction of aircraft temperature bias in the NCEP's GSI analaysis scheme. *Mon. Wea. Rev.,* **143,** 3774–3803.