

2B.3 REDUCING MOIST ADIABATIC CALCULATION COSTS USING LOOKUP TABLES

Nathan A. Dahl*

University of Oklahoma/CIMMS and NOAA/NWS/Storm Prediction Center, Norman, OK

1 INTRODUCTION

Calculating parcel temperature during undiluted pseudoadiabatic ascent is a significant source of computational cost in atmospheric model post-processing. Because there is no simple analytical solution for integrating the pseudoadiabatic lapse rate with respect to height, various empirical and numerical methods for obtaining sufficiently accurate final parcel temperature estimates have been proposed. For example, the “Wobus method” currently used in the National Centers Skew-T/Hodograph Analysis and Research Program (Hart et al. 1999) employs a secant method (Gerald and Wheatley 1984) to obtain a converged solution of wet-bulb potential temperature from wet-bulb temperature and a numerically-fitted “Wobus function.”

Noting the potential for parcel temperature errors in excess of 1 K when using the Wobus method, Davies-Jones (2008, hereafter DJ08) described an improved algorithm in which an inversion of Eq. (39) in Bolton (1980) is obtained empirically; the resulting formula produces a final parcel wet-bulb temperature estimate that is accurate to within 0.34 K relative to the fully-converged solution to Bolton (1980) Eq. (39), and one iteration of an adjustment scheme reduces the relative error to less than 0.002 K.

Along with the improved accuracy, this approach is more efficient than the Wobus method.

However, the empirical approximations used in the methods described above still require complex calculations that substantially contribute to the runtime required for post-processing. Another option is to perform the costliest calculations offline over a range of relevant values for the initial parcel state (e.g., temperature, dewpoint, pressure p) and the final pressure and store the results in a table. This leaves the post-processing program with only the relatively simple and inexpensive task of reading the table and interpolating the tabulated values to obtain the final temperature for a given parcel.

In an early example of this approach, Prosser and Foster (1966) estimated parcel wet-bulb potential temperature θ_w by precalculating temperatures along the $\theta_w = 10^\circ \text{C}$, 20°C , and 30°C pseudoadiabats at several pressure levels. They then estimated the final parcel temperature through linear interpolation of the precalculated data using the parcel wet-bulb temperature at $p = 1000$ mb and the final parcel pressure. Though inexpensive, this method was quite inaccurate (due to the coarseness of the precalculated data as well as the simple empirical approach used to obtain it) and limited to surface-based parcels.

A more comprehensive and accurate lookup method is currently implemented in the Weather Research and Forecast (WRF)

* Corresponding author address: Dr. Nathan Dahl, Cooperative Institute for Mesoscale Meteorological Studies, 120 David L. Boren Blvd., Suite 2100, Norman, OK 73072. email: nathan.dahl@noaa.gov

model. A table (psadilookup.dat) listing temperature at 150 pressure levels along 150 separate pseudoadiabats specified by equivalent potential temperature (θ_e) is read in at initialization. The temperature of a saturated parcel at a given pressure level is estimated by calculating the parcel θ_e and then applying bilinear interpolation to the tabulated values.

However, the cost of calculating the parcel θ_e is not trivial. Furthermore, the θ_e and pressure coordinates of the table are spaced irregularly (stretched from 200 to 688 K for θ_e and 40 to 1100 mb for pressure), which complicates the lookup process somewhat. (It should also be noted that this table has not been well publicized; its existence only recently became known to the author.) A method is proposed herein which mitigates both of these difficulties, enabling faster calculation of saturated parcel temperatures with comparable accuracy.

2 METHOD

A saturated parcel's temperature upon rising or sinking to pressure level p_f (given herein in mb) is determined by its lifted condensation level (LCL) temperature and pressure (T_{LCL} and p_{LCL} , in degrees C and mb respectively). We proceed on the assumption that all of the (T_{LCL}, p_{LCL}) combinations of meteorological significance (in terms of convection in Earth's atmosphere) are contained within the domain of a skew-T log-P diagram (hereafter "skew-T"). Using the isotherms and isobars of a typical skew-T for reference, we employ the following function to map the parcel from (T_{LCL}, p_{LCL}) to "skew-T coordinates" denoted (x_{LCL}, p_{LCL}) :

$$x_{LCL} = T_{LCL} + 90 \left[\frac{\log(1050) - \log(p_{LCL})}{\log(1050) - \log(100)} \right]. \quad (1)$$

x_{LCL} corresponds to the value of the temperature along the x-axis of the skew-T (i.e. at 1050 mb) directly beneath the T_{LCL} of the parcel. Using this coordinate as a basis for the table reduces the need to tabulate extreme values that have no real meteorological significance; for example, using temperature or potential temperature as a basis instead would require the inclusion of parcels well outside the bounds of actual observations in order to cover the full domain of a typical skew-T.

For the bounds of the table, it should be noted that $x_{LCL} \geq 40$ C is highly unlikely at present (although, with continuing climate change, this may not always be the case). The highest heat index currently on record was evaluated from observations taken on 8 July 2003 in Dharhan, Saudi Arabia; given the surface temperature of 42 C and the surface dewpoint of 35 C and assuming a surface pressure of roughly 1000 mb gives $T_{LCL} \approx 33$ C, $p_{LCL} \approx 900$ mb, and $x_{LCL} \approx 39$ C. Furthermore, while $x_{LCL} < -50$ C is possible in the upper levels of the atmosphere (or at lower levels in extremely cold climates), substantial vertical motion is not expected in such conditions. Therefore, the lookup table currently spans $-50 \text{ C} \leq x_{LCL} \leq 40 \text{ C}$ and $50 \text{ mb} \leq p_{LCL} \leq 1050 \text{ mb}$.

The method is based on a three-dimensional table containing "true" values of final wet-bulb parcel temperature $T_f(x_{LCL}, p_{LCL}, p_f)$ recorded at regular intervals Δx_{LCL} , Δp_{LCL} , and Δp_f . The "true" temperature at each location in the table is determined using Eq. (2.3) in DJ08, based on Eq. (39) in Bolton (1980):

$$f(T_W, \pi) = \left(\frac{C}{T_W}\right)^\lambda \left[1 - \frac{e_s(T_W)}{p_o \pi^\lambda}\right] \exp[-\lambda G(T_W, \pi)]$$

$$= \left(\frac{C}{T_E}\right)^\lambda \quad (2)$$

where

$$G(T_W, \pi) = \left(\frac{k_0}{T_W} - k_1\right) [r_s(T_W, \pi) + k r_s^2(T_W, \pi)]. \quad (3)$$

T_W is the final wetbulb temperature of the parcel in K (i.e. $T_W = T_f + C$ where $C = 273.15$ K), π is the final pressure p_f nondimensionalized by the base pressure $p_o = 1000$ mb, $\lambda = 3.5038$, e_s is the saturation vapor pressure in mb, T_E is the equivalent temperature of the parcel in K, $k_0 = 3036$ K, $k_1 = 1.78$, $k_2 = 0.448$, and r_s is the saturation mixing ratio. The solution is obtained by starting with a first guess τ_0 (obtained from DJ08 equations 4.8 through 4.11) and applying Newton's method until convergence (defined here as $|\tau_{n+1} - \tau_n| < 10^{-4}$ K):

$$\tau_{n+1} = \tau_n - \frac{f(\tau_n, \pi) - \left(\frac{C}{T_E}\right)^\lambda}{f'(\tau_n, \pi)}. \quad (4)$$

DJ08 derived explicit formulas for the first and second derivatives of f in order to speed up convergence. However, the speed of convergence is not a concern here because the solutions are calculated in advance (rather than at runtime) in this method. Therefore, for simplicity, we use centered differences to estimate f' in (4):

$$f'(\tau_n, \pi) \approx \frac{f(\tau_n + 0.01, \pi) - f(\tau_n - 0.01, \pi)}{0.02}. \quad (5)$$

The final temperature corresponding to the i th value of x_{LCL} , the j th value of p_{LCL} , and the k th value of p_f is obtained in this manner and recorded in the (i, j, k) location of the table as $T_{i,j,k} \equiv \tau_{n+1} - C$. As with the method used in WRF, the entire table is

populated in advance. At runtime, the table is read in; thereafter, any time a saturated parcel is evaluated, the parcel's (x_{LCL}, p_{LCL}, p_f) coordinates are obtained and the tabulated values are interpolated to those coordinates to estimate T_f .

3 DATA

During the 2019 NOAA Hazardous Weather Testbed Spring Forecast Experiment (SFE; Clark et al. 2019), participants were asked to subjectively rate the previous day's products from several real-time mesoscale analysis systems: the 40-km surface objective analysis (sfcOA) produced operationally by the Storm Prediction Center (SPC); the 15-km High Resolution Rapid Refresh Ensemble (Dowell et al. 2016) mean provided by the Earth Systems Research Laboratory (ESRL); the 3-km Warn-on Forecast System (WoFS; Skinner et al. 2018) ensemble mean, provided by the National Severe Storms Laboratory; and the 3-km 3-D Real-Time Mesoscale Analysis (3D-RTMA), developed jointly within NOAA by the NCEP Environmental Modeling Center and the ESRL Global Systems Division. The 3D-RTMA output was both presented on its native grid and coarsened to the 40-km sfcOA grid. All products were post-processed using the 3-D Analysis System (3DAS) developed by SPC. Participants overwhelmingly preferred the 3D-RTMA products. (Fig. 1)

The 3D-RTMA consists of various observations (surface station, radiosonde, radar, aircraft, satellite, lightning, cloud/hydrometeor, etc.) assimilated through Gridpoint Statistical Interpolation (Wu et al. 2002) onto a background state from the High Resolution Rapid Refresh model. The analysis grid spans the continental United States (CONUS), yielding

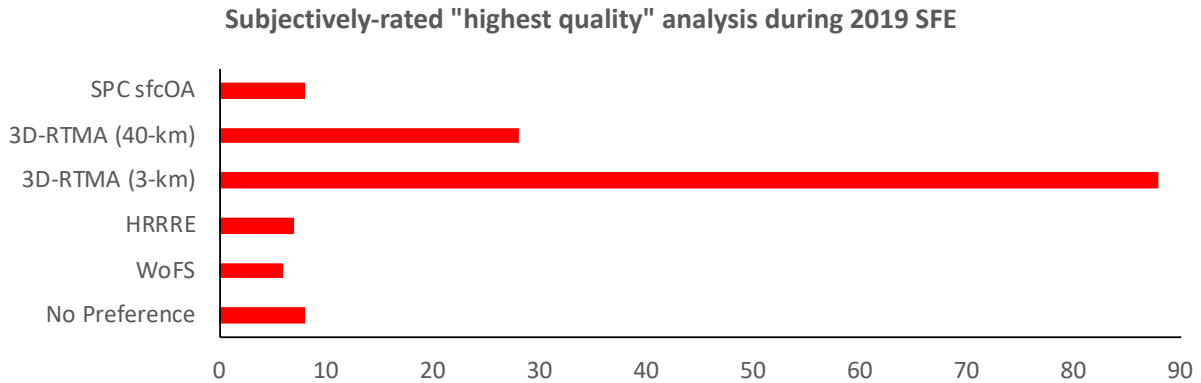


Figure 1 Results of survey asking 2019 NOAA HWT Spring Forecast Experiment participants to select which real-time mesoscale analysis system produced the highest-quality results. (from Clark et al. 2019)

nearly two million vertical profiles ($n_x = 1799$, $n_y = 1059$). Post-processing each vertical profile may require hundreds of pseudoadiabatic calculations; not only must parcels near the surface be lifted through ~ 50 vertical grid levels each, but several different parcels must be evaluated to obtain (for example) surface-based, mixed-layer, and most-unstable CAPE and CINH and to determine the base and top of effective inflow layers (Thompson et al. 2007).

It should be noted that previous tests (not shown) indicate that the increase in accuracy from an iteration of the DJ08 method over that of the first guess has very little impact on calculations of CAPE and CIN. It appears that this is due to the fact that the larger initial guess errors are concentrated near the discontinuities in DJ08 equations 4.8 through 4.11 (see Fig. 7 in DJ08) and thus would not have a substantial impact on parameters obtained by vertical integration through a deep layer of the atmosphere. Therefore, since the increased cost of a DJ08 iteration (compared to the cost of calculating the first guess alone) is not trivial, the DJ08 first guess (“DJFG”) is used as the baseline for

evaluating the speedup due to using a lookup table in 3DAS.

To assess the relative computational speed of the DJFG, WRF θ_e -based lookup table (“THeLU”), and “skew-T coordinates”-based lookup table (“SKEWLU”) methods within the context of general post-processing, four cases involving an “Enhanced Risk” or greater of severe weather were selected from the 2019 SFE. (Fig. 2) Three versions of 3DAS were prepared, one for each of the different pseudoadiabatic calculation methods. Each version was then run on the 3D-RTMA CONUS output for each case.

In previous tests (not shown), cases in which unstable conditions prevailed over a larger fraction of the domain were slower to process. To allow for this variation, and to test the effect of smaller domains in general, the 3DAS versions were also run on the corresponding WoFS domains ($n_x = 300$, $n_y = 300$) for each case (black boxes in Fig. 2). These domains focus more on regions of greater risk and therefore are covered in greater proportion by regions of substantial instability.




Figure 2 Storm Prediction Center convective outlooks and relevant times for four 3DAS test cases from the 2019 NOAA HWT Spring Forecast Experiment. Block boxes denote the corresponding Warn-on Forecast System domains.

4 RESULTS

4.1 Accuracy of interpolated estimates

To determine the impact of table resolution on the accuracy of the interpolated SKEWLU temperatures, 6 sets of intervals Δx_{LCL} , Δp_{LCL} , and Δp_f (with corresponding array dimension sizes Nx_{LCL} , Np_{LCL} , and Np_f) were chosen for testing as shown in Table 1.

Furthermore, because the manner of interpolation is expected to affect the accuracy of the estimate, two tests were performed on each table, one using linear interpolation for all three dimensions (denoted “lin-T lin-P”) and one using linear interpolation for the x_{LCL} dimension and logarithmic interpolation for the p_{LCL} and p_f dimensions (denoted “lin-T log-P”). The interpolation formulas are as follows:

Table 1: Dimensions and spacings of “skew-T coordinate”-based lookup tables.

	R1	R2	R3	R4	R5	R6
Δx_{LCL} (C)	10	5	2.5	1	0.5	0.25
Δp_{LCL} (mb)	50	25	10	5	2.5	1
Δp_f (mb)	50	25	10	5	2.5	1
Nx_{LCL}	10	19	37	91	181	361
Np_{LCL}	21	41	101	201	401	1001
Np_f	21	41	101	201	401	1001

$$T_1^* = (1 - \delta_2)(1 - \delta_3)T_{i,j,k} + (1 - \delta_2)\delta_3T_{i,j,k+1} + \delta_2(1 - \delta_3)T_{i,j+1,k} + \delta_2\delta_3T_{i,j+1,k+1}$$

$$T_2^* = (1 - \delta_2)(1 - \delta_3)T_{i+1,j,k} + (1 - \delta_2)\delta_3T_{i+1,j,k+1} + \delta_2(1 - \delta_3)T_{i+1,j+1,k} + \delta_2\delta_3T_{i+1,j+1,k+1}$$

$$T_f'(x_{LCL}, p_{LCL}, p_f) = T_1^* + \delta_1(T_2^* - T_1^*) \quad (6)$$

where a prime here denotes an interpolated estimate. The indices i , j , and k are given by

$$i = \text{floor}[(x_{LCL} + 50)/\Delta x_{LCL}] + 1$$

$$j = \text{floor}[(1050 - p_{LCL})/\Delta p_{LCL}] + 1$$

$$k = \text{floor}[(1050 - p_f)/\Delta p_f] + 1 \quad (7)$$

Here, “floor” indicates rounding down to the nearest integer and all pressures are given in mb. For lin-T lin-P, the interpolation weights δ_1 , δ_2 , and δ_3 are given by

$$\delta_1 = \frac{x_{LCL} - [(i-1)\Delta x_{LCL} - 50]}{\Delta x_{LCL}}$$

$$\delta_2 = \frac{[1050 - (j-1)\Delta p_{LCL}] - p_{LCL}}{\Delta p_{LCL}}$$

$$\delta_3 = \frac{[1050 - (k-1)\Delta p_f] - p_f}{\Delta p_f} \quad (8)$$

For lin-T log-P, the weights δ_2 and δ_3 are given by

$$\delta_2 = \frac{\log\{[1050 - (j-1)\Delta p_{LCL}]/p_{LCL}\}}{\log\{[1050 - (j-1)\Delta p_{LCL}]/(1050 - j\Delta p_{LCL})\}}$$

$$\delta_3 = \frac{\log\{[1050 - (k-1)\Delta p_f]/p_f\}}{\log\{[1050 - (k-1)\Delta p_f]/(1050 - k\Delta p_f)\}} \quad (9)$$

To assess the accuracy of each of the tables described in Table 1, a “truth” array of (x_{LCL}, p_{LCL}) coordinates is specified over $-50 \text{ C} \leq x_{LCL} < 40 \text{ C}$ and $100 \text{ mb} < p_{LCL} \leq 1050 \text{ mb}$, spaced every 0.1 C for x_{LCL} and every 0.5 mb for p_{LCL} . For each

(x_{LCL}, p_{LCL}) in the truth array, the fully-converged solution from (2)-(5) (T_f) and the interpolated estimate from a given table (T_f') are obtained for every p_f from 1050 mb to 100 mb₁ at 0.5 mb intervals.

The maximum differences between the converged solutions and the interpolated estimates (i.e. $\max[|T_f - T_f'|]$) over the range of p_f for each (x_{LCL}, p_{LCL}) are plotted in Fig. 3 (lin-T lin-P) and Fig. 4 (lin-T log-P). In both cases, the maximum errors for tables R3 through R6 are less than that reported for DJFG (0.34 C), and the maximum error for R6 is similar to that obtained from one iteration of the DJ08 method (0.002 C). Using logarithmic interpolation for pressure produces maximum errors up to 20-30% lower than those obtained through linear interpolation; however, the magnitude of the improvement is small and does not justify the 10-20% slowdown that results from using seen when lin-T log-P instead of lin-T lin-P (not shown).

Evaluation of the maximum errors for the interpolated THeLU temperatures followed a similar procedure. The converged solution from (2)-(5) was obtained for θ_e values from 200 to 688 K in 0.1 K increments at pressure levels from 1050 mb to 100 mb in 1 mb increments. This was then compared to the temperatures obtained by bilinear interpolation from the table. The result is shown in Fig. 5; the maximum error is approximately 0.01 C, putting the accuracy of this method on par with that of the SKEWTLU lin-T lin-P method using the R5 table. In light of these results, trilinear interpolation from the R5 table was selected for testing in 3DAS along with the DJFG and THeLU methods.

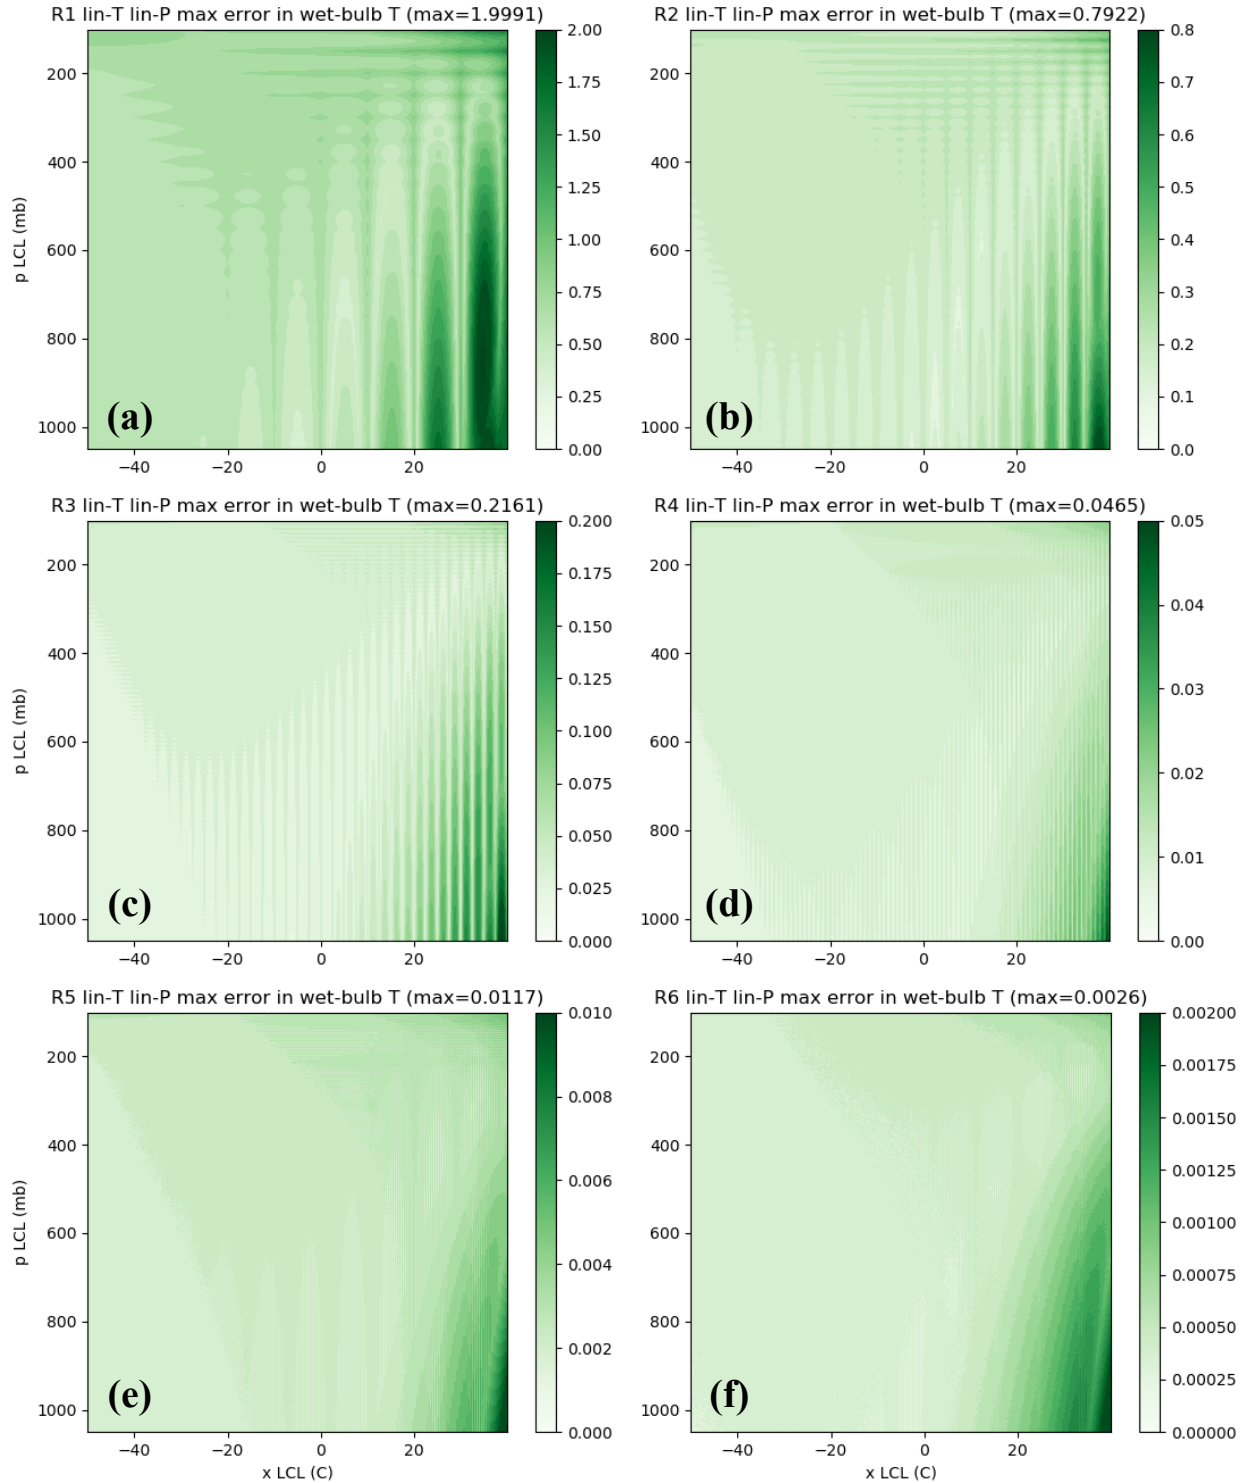


Figure 3 Maximum error (C) in the wet-bulb temperature estimate for a parcel undergoing saturated ascent or descent as a function of x_{LCL} and p_{LCL} , obtained by linear interpolation from lookup table “R1” (a), “R2” (b), “R3” (c), “R4” (d), “R5” (e), or “R6” (f). Lookup table characteristics are described in Table 1. Note the change in color bar range between panels.




Figure 4 Same as Fig. 3, but for estimates obtained using linear interpolation for x_{LCL} and logarithmic interpolation for p_{LCL} and p_f .

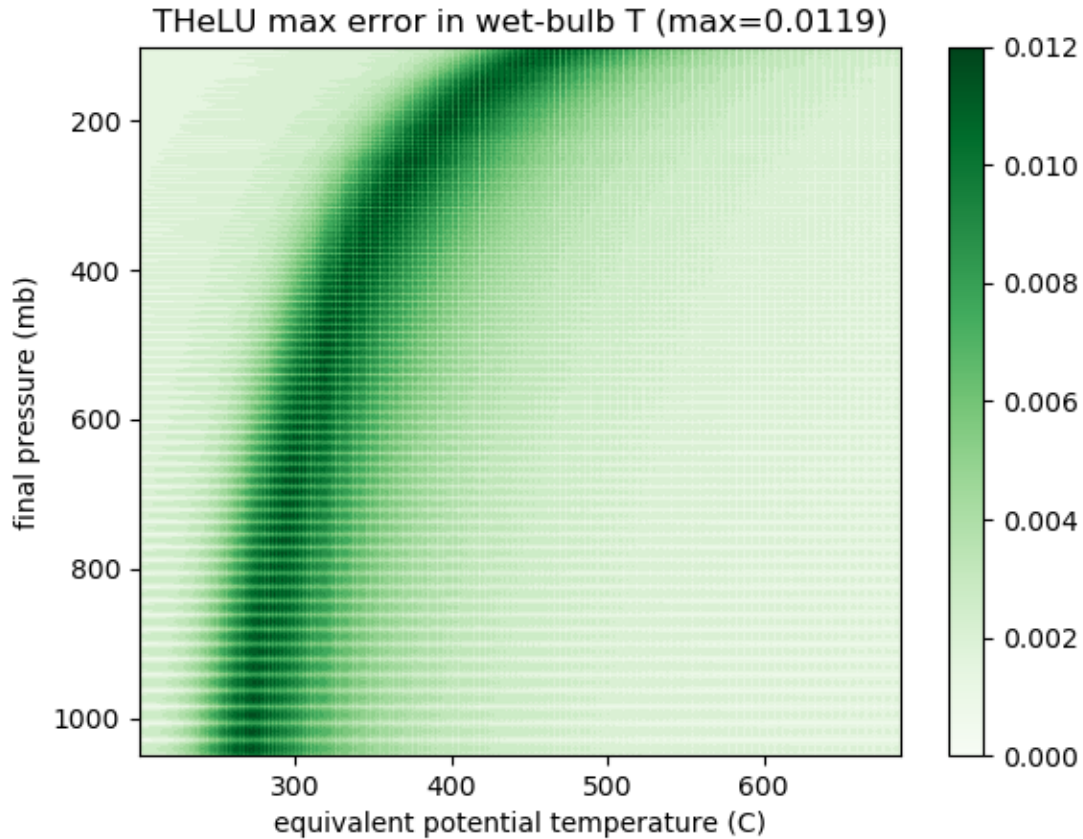


Figure 5 Same as Fig. 3, but for estimates obtained through bilinear interpolation of values in the θ_e -based table used in WRF.

4.2 Comparison of post-processing speeds

Table 2 lists the vertical profile processing rates and total runtimes for each of the methods applied to the 3D-RTMA fields. The SKEWTLU method is fastest by a considerable margin, consistently increasing the profile processing rate by approximately 50% relative to the TheLU method and 75% relative to the DJFG method. This amounts to a 33% reduction in total runtime relative to the DJFG method (compared with a 10% reduction for the TheLU method).

As described previously, processing the WoFS domains (see Fig. 2) is much slower due to a greater prevalence of moderately to

highly unstable profiles; as shown in Table 3, processing rates are generally reduced by over 50% relative to those obtained for the 3D-RTMA domains. On the other hand, the relative benefit of using a lookup table for pseudoadiabatic calculations is also greater; the improvement in processing rate is consistently around 95% for the SKEWTLU method (30% for the TheLU method), yielding a reduction in post-processing time of 35-40% (15-20% for the TheLU method).

It should also be noted that all three methods give very similar results for CAPE and CINH calculations. For example, the maximum analyzed CAPE is in the 4000-6000 J/kg range for each of the test cases, and the

Table 2: Profile processing speeds and total processing times for the Davies-Jones (2008) first guess (DJFG), θ_E -based lookup table (THeLU), and “Skew-T coordinate” lookup table (SKEWTLU) methods applied to 3D-RTMA fields for four test cases. Improvements relative to the DJFG method are shown in parentheses.

	processing speed (profiles/sec)			total runtime (sec)		
	DJFG	THeLU	SKEWTLU	DJFG	THeLU	SKEWTLU
30 Apr 2019	4805	5648 (18%)	8468 (76%)	500	451 (10%)	336 (33%)
01 May 2019	5019	5888 (17%)	8844 (76%)	490	439 (10%)	328 (33%)
06 May 2019	5054	5889 (17%)	8776 (74%)	485	439 (9%)	328 (32%)
07 May 2019	4981	5773 (16%)	8620 (73%)	494	445 (9%)	331 (33%)

Table 3: Same as Table 1, but for WoFS domains outlined in Fig. 2.

	processing speed (profiles/sec)			total runtime (sec)		
	DJFG	THeLU	SKEWTLU	DJFG	THeLU	SKEWTLU
30 Apr 2019	2093	2709 (29%)	4100 (96%)	56	46 (18%)	35 (38%)
01 May 2019	1957	2517 (29%)	3761 (92%)	59	49 (17%)	36 (39%)
06 May 2019	1874	2442 (30%)	3606 (92%)	61	50 (18%)	37 (39%)
07 May 2019	2021	2629 (30%)	3934 (95%)	57	47 (18%)	36 (37%)

maximum difference between methods is never above 30 J/kg. This is true even though the maximum error in parcel temperature for the DJFG method is 0.34 K, compared with 0.01 K for the other two methods. This suggests that using a much coarser lookup table (e.g. at least down to the “R3” specification; see Table 1 and Fig. 3) may be adequate for post-processing.

5 DISCUSSION

The SKEWTLU method can determine saturated parcel temperatures much more rapidly than previous methods can, without sacrificing accuracy. While the current work only examines its use in post-processing, there may be applications within forecast models as well, e.g., for pseudoadiabatic calculations when the Betts-Miller-Janjic

parameterization scheme (Betts and Miller 1986; Janjic 1994) is invoked.

As a caveat, the practical benefit of implementing this method is still unclear. For example, the tests in this study were only performed for the subset of mesoanalysis parameters used in the 2019 SFE. Roughly six times as many parameters are included in the full SPC mesoanalysis, most of which do not require any pseudoadiabatic calculation, and therefore the relative reduction in computational time and cost due to implementing SKEWTLU in the full mesoanalysis is expected to be somewhat less than what is shown here. Even so, it appears reasonable to expect such an implementation to reduce the required post-processing time by several minutes, with obvious benefits for forecasters.

ACKNOWLEDGEMENTS

This work is an adaptation of mesonalysis code provided by John Hart. The author also benefitted from collaboration with Greg Blumberg and Kimberly Hoogewind and assistance from the Storm Prediction Center Science Support Branch (Patrick Marsh, Israel Jirak, and Andy Dean) and Information Technology staff (Jay Liang and Gregory Grosshans). This extended abstract was prepared by Nathan Dahl with funding provided by NOAA/Office of Oceanic and Atmospheric Research under NOAA-University of Oklahoma Cooperative Agreement #NA16OAR4320115, U.S. Department of Commerce. The statements, findings, conclusions, and recommendations are those of the authors and do not necessarily reflect the views of NOAA or the U.S. Department of Commerce.

REFERENCES

- Betts, A. K., and M. J. Miller, 1986: A new convective adjustment scheme. Part II: Single column tests using GATE wave, BOMEX, ATEX, and arctic air-mass data sets. *Quart. J. Roy. Meteor. Soc.*, **112**, 693-709.
- Bolton, D., 1980: The computation of equivalent potential temperature. *Mon. Wea. Rev.*, **108**, 1046-1053.
- Clark, A., and Coauthors, 2019: Spring Forecast Experiment 2019: Program Overview and Operations Plan. 39 pp., https://hwt.nssl.noaa.gov/sfe/2019/docs/HWT_SFE2019_operations_plan.pdf
- Davies-Jones, R., 2008: An efficient and accurate method for computing the wet-bulb temperature along pseudoadiabats. *Mon. Wea. Rev.*, **136**, 2764-2785.
- Dowell, D., and Coauthors, 2016: Development of a High-Resolution Rapid Refresh Ensemble (HRRRE) for severe weather forecasting. *28th Conf. on Severe Local Storms*, Portland, OR, Amer Meteor. Soc., 8B.2, <https://ams.confex.com/ams/28SLS/webprogram/Paper301555.html>.
- Gerald, C., and P. Wheatley, 1984: *Applied Numerical Analysis*. 3rd ed. Addison-Wesley, 576 pp.
- Hart, J. A., J. Whistler, R. Lindsay, and M. Kay, 1999: NSHARP, version 3.10. Storm Prediction Center, National Centers for Environmental Prediction, Norman, OK, 33 pp.
- Janjic, Z. I., 1994: The step-mountain Eta coordinate model: Further developments of the convection, viscous sublayer, and turbulence closure schemes. *Mon. Wea. Rev.*, **122**, 927-945.
- Prosser, N. E., and D. S. Foster, 1966: Upper air sounding analysis by use of an electronic computer. *J. Appl. Meteor.*, **5**, 296-300.
- Skinner, P. S., and Coauthors, 2018: Object-based verification of a prototype warn-on-forecast system. *Wea. Forecasting*, **33**, 1225-1250.
- Thompson, R. L., C. Mead, and R. Edwards, 2007: Effective storm-relative helicity and bulk shear in supercell thunderstorm environments. *Wea. Forecasting*, **22**, 102-115.
- Wu, W.-S., D. F. Parrish, and R. J. Purser, 2002: Three-dimensional variational analysis with spatially inhomogeneous covariances. *Mon. Wea. Rev.*, **130**, 2905-2916.