



Truly Conserving with Conservative Remapping Methods

Karl E. Taylor

Program for Climate Model Diagnosis and Intercomparison, Lawrence Livermore National Laboratory, Livermore, CA 94550, USA

Correspondence: Karl E. Taylor (taylor13@llnl.gov)

Abstract.

Conservative mapping of data from one horizontal grid to another should preserve certain integral or mean properties of the original data. This may be essential in some model applications, including ensuring realistic exchange of energy and mass between coupled model components. It can also be essential for certain types of analysis, such as evaluating how far a system is from an equilibrium state. For some common grids, existing remapping algorithms may fail to perfectly represent the shapes and sizes of grid cells, which leads to errors in the remapped fields. A procedure is presented here that enables users to rely on the mapping weights generated by remapping algorithms but corrects for their deficiencies. With this procedure, for a given pair of source and destination grids, a single set of remapping weights can be applied to remap any variable, including those with grid cells that are partially or fully masked.

1 Introduction

When analyzing climate data from different sources, it is often necessary, as an initial step, to map the data to a common grid, a procedure commonly referred to as remapping or “regridding” the data. For some purposes it is essential when remapping the data that the global mean (or, alternatively, the global integral) of the field be preserved. Conservative remapping algorithms are meant to guarantee this. In practice, remapping occurs in two steps: 1) given a source and destination grid, mapping “weights” are computed, and then 2) a sparse matrix multiplication of the source data by the weights yields the values of the field on the destination grid. Our focus here is on the second step: given the weights needed for remapping conservatively, we provide guidance on how they should be applied. Appendix A lists some remapping packages that can be used to generate weights (i.e., to execute step 1). It should be noted that many (perhaps all) of these packages slightly misrepresent the true grid cell shapes for certain common grids, so care must be taken to correct for this. Moreover, most packages provide inadequate guidance on how to handle fields that are partially masked or for three dimensional fields how to account for variations in the thickness of individual layers. The main purpose here is to clearly explain how to compensate for a remapping algorithm’s inaccurate representation of grid cell shapes and to account for missing or partially masked data when that is necessary.

The objective in remapping conservatively is to preserve certain physically important characteristics of the climate system. For a climate in global thermodynamic equilibrium, for example, the mean net flux of energy at the top of the atmosphere is zero. In properly formulated models run with all externally-imposed conditions unchanging over time, the net flux at the top of



the atmosphere will indeed approach and fluctuate about zero as the system approaches equilibrium. When the fluxes from a simulation of this kind are mapped to a different grid, we would like to preserve this important characteristic of the simulation. This can only be done if the remapping algorithm is conservative.

30 The fundamental relationship that must be satisfied to preserve the integral of a field over the global area is:

$$r_s^2 \sum_i F_{si} f_{si} A_{si} = r_d^2 \sum_j F_{dj} f_{dj} A_{dj}, \quad (1)$$

where F is the quantity that is mapped from a source grid to a destination grid, A is the grid cell area (expressed as a solid angle and globally spanning 4π steradians), f is the grid cell fraction where F is defined (i.e., the “unmasked fraction” of a cell), and the subscripts, s and d , distinguish among the source and destination cells, respectively. The i and j indices apply to source and destination cells, respectively, and the sums are over the entire domain. In certain modeling applications, the radius of the earth, r , for the source and destination grids may differ slightly, and this is accounted for by including the squares of r_s and r_d in (1). If there is a mismatch in radii, the values of F_d must be scaled so that the integral is preserved. The sums are over all grid cells in the domain. For source cells that are completely masked or where data are undefined or “missing”, $f_s = 0$. For the most common case when the value in a source cell is representative of the entire cell extent and when there are no missing data, $f_s = 1$ for all grid cells. Note that when remapping conservatively, if $f_s = 1$ for all source cells, then we require that $f_d = 1$ for all destination cells.

A different relationship must be satisfied to preserve a field’s global mean (denoted by an overbar):

$$\overline{F_s} = \frac{\sum_i F_{si} f_{si} A_{si}}{\sum_i f_{si} A_{si}} = \frac{\sum_j F_{dj} f_{dj} A_{dj}}{\sum_j f_{dj} A_{dj}} = \overline{F_d}. \quad (2)$$

In the next section, we introduce the formulas that apply to remapping data when $f_s = f_d = 1$ everywhere (i.e., when there is no partial or complete masking of any cells). This is followed in section 3 by a description of the more general procedure when a field might be partially masked. In section 4 we provide recipes that should be followed in remapping different types of two-dimensional and three-dimensional fields. That is followed by a brief discussion of how to preserve certain properties when interpolating data in the vertical. In the summary, we include some discussion of when conservative remapping may be inappropriate (or at least inadvisable).

50 2 Remapping without masking

Consider first the simple case in which $f_s = f_d = 1$ for all grid cells. Weights, w_{ij} , that preserve the global mean will be used to calculate F_{dj} through a matrix multiplication:

$$F_{dj} = \sum_i w_{ij} F_{si}. \quad (3)$$

These weights would need to be scaled by r_s^2/r_d^2 , to preserve the global integral, rather than the global mean. Initially, we shall suppose that it is the global mean that should be preserved. At the end of the section 3, we show that the destination field that preserves the global mean can be simply scaled to preserve the integral.



Conceptually, the remapping weights are determined by overlaying the destination grid onto the source grid and calculating what fraction of each source cell overlaps each destination cell. These fractional contributions, here denoted ω , are then multiplied by the ratio of the source cell area to the destination cell area to yield:

$$60 \quad w_{ij} = \omega_{ij} A_{si} / A_{dj}. \quad (4)$$

Each distinct portion of a source cell must contribute to one and only one destination cell and all portions must contribute to some destination cell. This implies that within any source cell i , the fractional contributions ω_{ij} must sum to unity:

$$\sum_j \omega_{ij} = 1 \quad \text{for all } i. \quad (5)$$

This identity can be used to prove that for the case $f_s = f_d = 1$, (3) preserves the global mean. The proof is obtained by substituting (4) into (3), then substituting the result into (2), reversing the order that the summations are performed, and making use of (5). Importantly, a consistent set of areas must be used in evaluating (1), (2), and (5).

Noting that the fractional areas of the source cells contributing to a destination cell must add up to the area of the destination cell, we obtain a second useful identity:

$$\sum_i \omega_{ij} A_{si} = A_{dj} \quad \text{for all } j. \quad (6)$$

70 As discussed below, this identity holds only when the fractional contributions, ω , and cell areas, A , are consistently defined (i.e., based on the same cell shapes).

The most difficult step in remapping conservatively is to devise an algorithm to calculate the fractional contributions, ω_{ij} . Even when efficiently done, the calculation of ω_{ij} dominates the remapping execution time. In computing ω_{ij} , existing general remapping algorithms (see Appendix A) require the locations of grid cell vertices be specified (for both source and destination grids), and the remapping algorithm must make assumptions as to how the cell vertices are connected. The usual assumption is that the cell edges run along great circles even though for a cartesian longitude by latitude grid, the latitude cell bounds follow latitude circles, not great circles. For commonly used grids of this kind, the ω_{ij} calculated by the remapping algorithm are, in general, only approximations of the true fractional contributions because the assumed cell shapes are incorrect. This means the destination grid cell values will be somewhat misrepresented, and the remapping errors will generally exceed the formally derived error estimates.

80 After computing the fractional contributions, ω_{ij} , remapping codes can generate the remapping weights using (4). This ensures that

$$\sum_i F_{si} A_{si} = \sum_j F_{dj} A_{dj}, \quad (7)$$

where the areas are calculated based on the algorithm's assumed cell shapes. These areas may differ somewhat from the true cell areas. If they do, then in general

$$\sum_i F_{si} A_{si} \neq \sum_i F_{si} A_{si}^*, \quad (8)$$



where the asterisk distinguishes a true cell area (A^*) from an approximate cell area (A) generated by the remapping algorithm. The sum on the right hand side of the equation represents the true integral on the original source grid, but it is the sum on the left hand side that is preserved by (7).

90 In order to preserve the true global mean, some packages can, as an option, use true cell areas (provided by the user) in calculating the remapping weights,

$$w_{ij}^* = \omega_{ij} A_{si}^* / A_{dj}^*, \quad (9)$$

but it should be noted that in this case and in contrast with (6),

$$\sum_i \omega_{ij} A_{si}^* \neq A_{dj}^*. \quad (10)$$

95 The true fractional areas of the source cells contributing to a destination cell may not add up to the true area of the destination cell because the true areas may be inconsistent with the shapes of cells assumed in generating the fractional contributions, ω_{ij} .

Despite this apparent problem, some users may choose to calculate the destination values according to

$$F_{dj}^* = \sum_i w_{ij}^* F_{si}. \quad (11)$$

It follows that

$$100 \sum_i F_{si} A_{si}^* = \sum_j F_{dj}^* A_{dj}^*. \quad (12)$$

The sum on the left side of the equation represents the true global integral. Thus, when true areas are used in constructing the weights, the remapped field can preserve the true global mean.

Both alternatives for computing remapping weights, (4) or (9), rely on the same fractional contributions, ω_{ij} , which are based on approximate cell shapes constructed by the remapping algorithm. With the first option, the “approximate-area option”, cell areas are based on cells bounded by great circle segments, and the weights (w_{ij}) are defined by (4). In this case the areas are
 105 consistent with the assumed cell perimeters, so (6) holds. With the second option, the “true-area option”, the true grid cell areas must be supplied to the remapping algorithm, and the weights (w_{ij}^*) are defined by (9). In this case the areas may be inconsistent with the assumed cell perimeters used to generate the fractional contributions (ω_{ij}), so there may be a mismatch between the source areas and the destination areas, as noted in (10).

110 For the purpose of evaluating the relative merits of the two options, we now consider an example of a simple source grid and an idealized temperature field. For this example, suppose both the source and destination grids are spherical coordinate grids with the same latitude spacing (15°) but with the destination grid having half the longitudinal resolution of the source grid (60° vs. 30° spacing). Suppose the grids are aligned such that each destination cell completely contains exactly two source cells. The cell areas are given in Table 1. For the source cells nearest the pole, the true cell areas (with latitude cell bounds
 115 following latitude circles) is $0.0178 r^2$, whereas the approximate cell areas (assuming all cell sides follow great circles) is $0.0171 r^2$. Thus, there is a 4% error in approximating the cell area. In general these errors increase toward the poles and as grid



cell longitudinal width increases. For the destination grid, with twice the cell widths, the error quadruples to 17%, whereas halving the cell width shrinks the error to 1%. It can be shown that for small longitudinal cell widths, the fractional error in the approximate cell areas is proportional to the square of the cell widths.

120 In this example, correctly remapping a source field to the destination grid is trivial since each destination value is determined solely by the contributions from the two source cells that alone occupy it. Consider a temperature that varies linearly with latitude and is independent of longitude. Then, the destination field is identical to the source field but with half the longitudinal resolution. The temperature dependence on latitude for the case considered is given by the black line labeled “source grid (truth)” in Figure 1. Of course when correctly remapped, the destination values should in this case be the same as the source
125 values and also would lie on the black line.

If, however, we remap this temperature field based on an algorithm that assumes cell boundaries are defined by great circle segments when computing cell areas (the approximate-area option), the destination values will lie on the dashed brown curve in Figure 1. And if the algorithm relies on the true cell areas when computing weights (the true-area option), the discrepancy is much larger, with the resulting destination values given by the dashed blue curve in Figure 1. Neither option correctly remaps
130 the field, but the approximate-area option appears to be far superior to the true-area option.

Under each option, a global integral can be preserved, according to (12) and (7), but only when true areas are used and the destination field is obtained using (11) can we be certain to preserve the global integral as calculated on the original grid. Since the primary purpose in applying a conservative remapping scheme is to preserve the *true* global mean (i.e., the mean calculated with true areas), the approximate-area option would seem to be unacceptable. In Figure 1, although the destination
135 values shown for the approximate-area option appear to nearly coincide with the source values (“truth”), they are in fact systematically underestimated and lead to a global mean temperature 0.45 K cooler than the true mean.

On the other hand, it would seem equally unsatisfactory to adopt the true-area option (dashed blue curve of Figure 1), which produces remapped values at some latitudes differing by more than 30 K from the true values. The true-area option does indeed preserve the true global mean, but the pattern of the destination field can hardly be considered a good representation of the
140 source field. Thus, for different reasons, both options might be considered unacceptable.

It is interesting and somewhat disconcerting to note that with the true-area option, the results of remapping depend on the units used to express the temperature. The dashed blue curve in Figure 1 shows results when temperature is expressed in kelvins (K). The figure also shows that converting the temperature to degrees Celsius (°C), remapping that field, and then converting it back to kelvins results in considerably smaller errors (crosses in Figure 1). It can be shown that in general, the errors in
145 destination values, when computed using the true-area option, are approximately proportional to the magnitude of the values themselves. Since, on average, we can reduce the mean of the absolute values of a field by removing the global mean before remapping, we can use this strategy to reduce the errors. Our example of converting the temperature units from K to °C was an approximate application of this strategy which reduced the error because 0° C is much closer to the mean than 0° K. If we were to adopt the more general approach of removing the global mean before remapping, we arrive at the following formula:

$$150 \quad F_{dj}^* = \overline{F}_s + \sum_i w_{ij}^* (F_{si} - \overline{F}_s), \quad (13)$$



where here and in what follows an overbar indicates a global mean that must be computed using the true areas, not the approximate areas that remapping algorithms might generate. This variant of the true-area option will subsequently be referred to as the “true-area (centered)” approach, to distinguish it from the “true-area (uncentered)” approach which relies on (11).

155 An objection to using (13) is that a change in F_s anywhere in the domain will change the mean and thereby impact the destination values everywhere in the domain. One would not expect a remote change in a field to affect a local destination value, so (13) would also seem to be less than ideal. It should be noted that if weights generated with the approximate-area option were used in (13), the resulting destination field would be identical to that obtained with (3). This is because for these weights, (6) holds.

160 Yet another shortcoming of the true-area option is that its application to a spatially uniform source field can result in a destination field with non-zero spatial variance, which is obviously unrealistic. Consider, for example, a source field that everywhere has the value 1. For the grid defined earlier (see Table 1), application of (11) results in a destination field with the same mean (equal to 1), but with non-zero, area-weighted variance equal to 0.064, and a maximum absolute deviation from the true value of 0.13. These unrealistic variations may in some applications be unacceptable. Algorithms that maintain the uniformity of an originally constant field are said to be “consistent” (e.g., Ullrich and Taylor (2015)), so the true-area option
165 might be described as “not consistent”.

For the true-area option, use of (11) or (13) can sometimes result in a destination field with nonphysical values. Consider, for example, a possible result of mapping to a destination grid the ice-free fraction (i.e., the fraction of a grid cell area that is ice-free). As in the first example above, suppose ice-free fraction is independent of longitude and a linear function of latitude, varying from 0 in the polar-most latitude band to 1 in the latitude band adjacent to the equator. Application of (11)
170 results in a value of 1.06 for the latitude band nearest the equator, and application of (13) results in a value of 1.01. Clearly, the remapping algorithm in both cases yields a nonphysical result with the ice-free fraction exceeding 1 in the tropics. In contrast, the approximate-area option generates destination values that never exceed the maximum or minimum source values. The approximate-area approach is said to be a monotone method (e.g., Ullrich and Taylor (2015)), whereas the true-area approach is not.

175 Given the shortcomings of both the centered and uncentered variants of the true-area option, we reconsider the approximate-area option, which relies on the remapping algorithm to construct cell shapes and areas assuming perimeters coincide with great circle segments. The fundamental problem with this approach, as expressed by (8), is that the true global mean (as calculated on the source grid using true areas) is not generally preserved. In the fields we have examined, we have found that the difference between the true mean and the mean of the field remapped using the approximate area approach is quite small.
180 It seems reasonable, therefore, to simply adjust all values in the field by a uniform amount to correct for the small mismatch in means. For the case considered in Figure 1, 0.45 K can be added to each of the destination grid cell values obtained with the approximate-area option. This straight-forward adjustment eliminates the flaw in the approximate-area option that led us to discard it originally. This “global adjustment” to the destination field means that local destination field values will be influenced by remote field values. This is undesirable, but as noted earlier, the true-area (centered) approach is similarly impacted. On the
185 other hand, a virtue of the approximate-area option is that unlike the true-area option, a change in units (from, for example,



kelvins to °C) does not affect the accuracy of the result. In addition, a source field that has no variations (i.e., is everywhere the same) will map to a destination field that is also constant. Both of these results follow because when applying the approximate-area option, (6) holds.

190 Since neither of the cell-area choices offered by remapping codes is without shortcomings, it is worth further examining the characteristics of their errors to determine which approach results in the more realistic representation of the original field. For the temperature field considered earlier, Figure 2 shows the remapping error when approximate cell areas are used (with and without a global mean correction) and when the true cell areas are used with global mean removed and then reapplied following (13), labeled the “true area (centered) option” in the figure. By design, the correction to preserve the global mean under the approximate-area option simply offsets the curve by the same amount everywhere. With this correction and ignoring
195 for the moment the dotted curve, the “approximate-area (corrected uniformly) option” (as it will be identified subsequently) has the smallest root-mean-squared error (RMS error, calculated with grid cell area weighting). The RMS errors are: 0.15 K for the approximate-area (corrected uniformly) option, 0.48 K for the approximate-area (uncorrected) option, and 0.63 K for the true-area option (centered). In this example, although the correct global mean is preserved under two of the methods, the approximate-area (corrected uniformly) option leads to an RMS error a factor of 4 smaller than that resulting from the true-area
200 (centered) option.

In Figure 3 we consider for a spherical coordinate (cartesian longitude by latitude) grid how the remapping errors depend on grid cell size. The source data were taken from Figure 1 with values independent of longitude location and 15° latitude spacing. Source grids with longitudinal resolution from 0.5° to 30° were considered, and in each case the destination grid had half the longitudinal resolution of the source. As discussed above, the errors arise because of inaccuracies in the representation of grid
205 cell shapes by the remapping algorithm (which assumes the cell latitude bounds follow great circles, not latitude circles). It is not surprising then that the finer the resolution, the smaller the errors (because for small grid cells, the great circles deviate very little from latitude circles). Figure 3 shows that under all options the RMS error is proportional to the square of the grid cell’s longitude width and that for any given resolution, the approximate-area option error is about 1/4 the size of the true-area (centered) option error and more than two orders of magnitude smaller than the error in uncentered true-area option. From
210 another perspective, compared to the true-area (centered) option, the approximate-area option can with equal accuracy handle grids that are twice as coarse.

As noted earlier, when there are some physical limits on a variable (e.g., a fraction confined to the interval 0 to 1), remapping algorithms may not respect those limits. Although with the approximate-area option, the remapping step ensures that all destination values will be within the maximum and minimum values of the source values, the correction to the mean required
215 when applying that option can sometimes push values outside the limits. This issue can be addressed with a refined correction, which will be described in a part of the next section.



3 Remapping of partially or fully masked cells

We now consider the more general procedure for conservative remapping when there might be undefined elements in the source array (e.g., missing or masked elements) or when grid cell values might be defined for only a fractional portion of the source cell (for example only over the land portion of a cell). For this purpose, we will adopt and generalize the form of the approximate-area option because, as discussed above, it was found to be more accurate than the true-area options and because with this option we can simplify some subsequent formulas using (6). Moreover, where the field is constant on the source grid, the approximate-area option, unlike the true-area option, does not introduce unrealistic spatial variations in a region.

The key to handling data that may represent conditions on only a portion of each grid cell is to specify for each cell the “unmasked” fraction, and when remapping is performed, generate the appropriate destination unmasked fractions. Although sometimes the unmasked fractions are binary (either 0 or 1) and might be inferred from special bit strings assigned to any “missing” data, if the data are remapped, the unmasked fractions will in general no longer be binary, and thus information will be lost unless the unmasked fractions are carried as an additional field along with the data field. The key to general remapping then is to carefully account for the unmasked fractions and to ensure that they are consistently defined on the destination grid.

Generalizing (3) to account for fully or partially masked fields requires modification of the weights defined by (4). This is done by replacing in (4) the ratio of areas by the ratio of *unmasked* portions of the cell areas. Alternatively and equivalently, we can explicitly include the unmasked fractions in (3), which then becomes

$$\hat{F}_{dj} = \sum_i w_{ij} (f_{si} / \hat{f}_{dj}) F_{si}, \quad (14)$$

where the ‘hat’ atop F_{dj} indicates that this is a preliminary estimate of the destination value which might be corrected subsequently to preserve the global mean. Note that the source unmasked fractions, f_{si} , must be set to 0 wherever there are missing data. When this is done, missing data need not be treated specially because whatever value has been assigned to indicate the data value is missing, the value will invariably get multiplied by f_{si} , yielding 0 for the product. This ensures that missing values have no impact on the remapped fields.

For some applications, destination fractions may have been imposed as part of the definition of the destination grid. For the purposes of remapping a field, however, it is essential that the destination fractions in (14) be defined such that

$$\hat{f}_{dj} = \sum_i w_{ij} f_{si}. \quad (15)$$

The remapping formula (14) can then be written

$$\hat{F}_{dj} = \frac{\sum_i w_{ij} f_{si} F_{si}}{\sum_i w_{ij} f_{si}}. \quad (16)$$

Thus, the destination value is simply a weighted mean of the contributing grid cell values. This ensures that the destination value will not lie outside the maximum and minimum values of the contributing cells. This further implies that if all source cells contributing to some destination cell have the same value, the destination cell will also be assigned that value. Note that if in (16) the sum in the denominator is zero, then the destination value should be designated as missing.



As shown earlier, use of approximate areas in computing the weights in (16) does not in general preserve the true global mean of F . As in the simpler case, a global adjustment to the \hat{F}_{dj} values must be applied, but here we allow the correction to vary spatially,

$$F_{dj} = \hat{F}_{dj} - \frac{\gamma_j}{\bar{\gamma}} (\hat{F}_d - \bar{F}_s), \quad (17)$$

where the global mean of γ in the denominator above ensures $\bar{F}_d = \bar{F}_s$. With masking, the global mean quantities (indicated by overbars) must be computed with area weights proportional to only the *unmasked* area of grid cells, For example,

$$\bar{F}_d = \frac{\sum_j \hat{F}_{dj} f_{dj} A_{dj}^*}{\sum_j f_{dj} A_{dj}^*}. \quad (18)$$

Note that in this formula, the unmasked fraction, f_a , is not necessarily identical to the unmasked fraction, \hat{f}_a , which appears in (14). Sometimes, for example, the remapped data must conform to a destination grid with an imposed masked region. In that case, the already defined fractions, f_a , can be used in (18). This, however, could result in some destination field values calculated with (16) being masked. With those values no longer contributing to the destination field, the correction to the mean given by (17) must compensate, and this will alter the destination values, F_{dj} , globally, not just locally. It is therefore advisable to assign destination masked fractions consistent with (15) and avoid imposing externally defined destination masked fractions.

In (17), the correction coefficient, $\gamma_j/\bar{\gamma}$, can vary spatially, but in the simplest case is set to 1 for all j , which adjusts each cell value by the same amount ($\hat{F}_d - \bar{F}_s$). Sometimes this can lead to nonphysical results. Suppose, for example, that a positive definite quantity (such as the liquid water content of air) were mapped to a target grid using (16) and that the resulting global mean were greater than \bar{F}_s . In this case, any cell with $\hat{F}_{dj} = 0$ would, after a simple adjustment with $\gamma_j = 1$, become negative, which must be ruled out on physical grounds.

More generally, a uniform adjustment of the destination field may result in values that lie outside the range of source values. Returning to our earlier example, we see in Figure 2 that a uniform correction to the temperature field of 0.45 K results in a positive 0.1 K error in the equator-most cell, and the remapped temperature there is warmer than the warmest temperature found in the original field (290 K). Thus, the remapped field, which before correction was monotonic, is no longer so.

To remedy this undesirable consequence of a uniform correction, we suggest the correction coefficient, γ , be of the form

$$\gamma_j = (\hat{F}_{dj} - F_{\min})^\mu (F_{\max} - \hat{F}_{dj})^\mu, \quad (19)$$

with $F_{\min} = \min[F_{si}]$ and $F_{\max} = \max[F_{si}]$. The value of μ should be chosen just large enough to prevent in the destination field extremes beyond the source field limits. When μ is larger than necessary, the adjustments will unevenly be distributed among relatively few grid cells and may lead to a less accurate result. When the values of F_s span multiple orders of magnitude, the correction, if evenly distributed, would have a much larger relative impact on the smallest values. Consider, for example, atmospheric smoke concentration, which might be found in a limited region. Any correction to ensure the mean is conserved, should mainly be applied to the high concentration region. In this case, then, μ should be large enough that the low-level, background concentrations of smoke (all near 0) would be unaffected. Note that one could devise an even more versatile adjustment by using two different exponents for the two factors in (19)



280 As a rule in calculating γ , a relatively small value of μ , say 0.1 or 0.25, should be tried first, and then the destination values should be checked to confirm that none exceeds the extremes found in the source field. If this condition is not met, the value of μ should be incrementally increased and γ recalculated until all destination field values fall within the extremes of the source field. In a few test cases, we found that the smaller the value of γ , the smaller the overall remapping error.

In the example considered above, a correction with $\gamma = 0.25$ can be applied that eliminates the unrealistically high temperature near the equator which resulted from a *uniform* correction. The dotted curve in Figure 2 shows the result. In addition to
285 eliminating the artificial maximum, the correction slightly reduces the RMS error in the remapped field (see Figure 3), but this result does not carry over to all the resolutions considered. Larger values of γ would reduce the temperature correction near the extremes, which would be compensated by larger corrections to intermediate temperatures, as shown in Figure 4.

In this example involving coarse resolution grids (with longitude widths $\geq 30^\circ$), the uncorrected global mean of the remapped
290 field is 0.45 K cooler than the true global mean. For finer resolution grids, the magnitude of the correction is considerably smaller: when, for example, the same temperature field is mapped from a source grid with cell widths of 4 degrees longitude to a grid with widths of 8 degrees, the global mean is less than 0.01 K cooler than the true mean. When it is *not* essential to preserve the true mean, a user might choose to accept such a small error in global mean in order to avoid the correction procedure described above.

295 A final adjustment to F_{dj} may be needed if the objective is to preserve the global *integral* of a field, rather than the global *mean*. (To conserve energy in a coupled climate model, for example, it is the surface heat flux between the atmospheric component and the ocean component that might need to be mapped from one grid to another, and it is the total flux which must be preserved.) Comparison of (1) and (2) indicates that to preserve the integral, the values of F_{dj} obtained using (17), which preserve the mean, should be scaled by the ratio of the global unmasked source area to the unmasked destination area:

$$300 \quad c = \frac{r_s^2 \sum_i f_{si} A_{si}^*}{r_d^2 \sum_j f_{dj} A_{dj}^*}. \quad (20)$$

If the destination unmasked grid-cell fractions have been defined such that the global mean unmasked fraction is preserved, the sums in the numerator and denominator of (20) are the same, and the formula simplifies to $c = (r_s/r_d)^2$. When this is true and if $r_d = r_s$, the same destination field, without scaling, preserves both the mean and the integral. Note, however, if the destination unmasked fractions have been calculated with (15) and they have *not* been corrected to preserve the global mean
305 unmasked fraction, then c must be calculated with (20).

4 Recipes for regriding

Some conservative remapping packages (see Appendix A) have not been designed to handle the most general cases considered here (e.g., fields with missing values or grid cells that are partially masked). Those codes may nevertheless be relied on to provide weights defined by (4). For a given source and destination pair of grids, these weights can be calculated once, and then
310 used to remap any field from the source grid to the destination grid.

A step-by-step procedure for mapping variables conservatively is provided here.



1. Obtain from a remapping package the weights (w_{ij}) that apply when there is no masking or fractional weighting and that are generated based on the algorithm's approximate reconstruction of grid cell shapes.
2. Check that for all destination cells the weights satisfy (6), which with (4) can be rewritten:

315
$$\sum_i w_{ij} = 1 \quad \text{for all } j. \quad (21)$$

(It is prudent to include this step but not necessary if the weights returned by the remapping package are known to satisfy this requirement.)

3. Assign or calculate the source grid's unmasked fractions, f_{sj} .
 - If a source value is meant to represent conditions over the entire cell extent, set $f_{si} = 1$ for the cell.
 - 320 – If unmasked fractions have been assigned to source cells prior to remapping, the pre-assigned values should be assigned to f_s .
 - Wherever source cell data are missing, reset the unmasked fraction to 0 ($f_{si} = 0$).

4. Assign or calculate the destination grid's unmasked fractions, f_{dj} .
 - If unmasked fractions have been assigned to destination cells prior to initiating the remapping procedure, f_d should
 325 be set to the pre-assigned values. As noted in the previous section, however, when destination masked values are not defined by (15), the destination field, F_d , will be impacted everywhere, so when possible, avoid applying external destination masked fractions different from \hat{f}_d .
 - If unmasked fractions have not been pre-assigned, generate preliminary estimates of the fractions with (15). When
 necessary or desirable, correct the values of \hat{f} to preserve the global mean by applying formulas analogous to (17)
 330 and (19):

$$f_{dj} = \hat{f}_{dj} - \frac{\gamma_j}{\bar{\gamma}} (\hat{f}_d - \bar{f}_s) \quad (22)$$

and

$$\gamma_j = \hat{f}_{dj}^\mu (1 - \hat{f}_{dj})^\mu. \quad (23)$$

335 This step ensures that the same destination field will preserve both the global integral and mean under the conditions discussed following (20). If the values of \hat{f} are corrected to preserve the true global mean using (22), a uniform adjustment ($\gamma_j/\bar{\gamma} = 1$) cannot normally be applied as that could lead to some negative value of f or some values exceeding 1; rather, the adjustment should be made with γ defined by (19) and $\mu > 0$. The global mean quantities (\bar{f}_d and \bar{f}_s) in the first equation above must be calculated using the true areas (A_{dj}^* and A_{si}^*) and without any fractional weighting.



- 340 5. Use (16) to calculate the preliminary destination values, \hat{F}_{dj} . For any cell where the denominator in (16) is 0, designate the destination value as "missing".
6. Correct the destination values, \hat{F}_{dj} , to preserve the true global mean and obtain the final destination field, F_d . In the next two sub-steps, when calculating (17), global means of F_{si} and \hat{F}_{dj} should be weighted by $f_{si}A_{si}^*$ and $f_{dj}A_{dj}^*$, respectively.
- 345 – Initially, impose a uniform correction to all values by applying (17) with $\gamma_j = 1$ for all j . If none of the resulting F_{dj} values have been shifted outside the extremes of the source field, consider the correction acceptable; otherwise proceed.
- If necessary, next try applying (17) with γ defined by (19) and μ initially set to some relatively small value (e.g., 0.1 or 0.25). As before, if none of the resulting F_{dj} values has been shifted outside the extremes of the source field,
- 350 consider the correction acceptable; otherwise repeat this step with μ increased (say doubled). Repeat this step, as necessary, until the correction becomes acceptable.

Note that no correction of \hat{F}_{dj} is needed if two conditions are met: 1) the remapping algorithm correctly represents the grid cell shapes (in which case the fractional contributions, ω , and cell areas will both have been correctly determined), and 2) the unmasked fractions on the destination grid are defined by (15).

355 In what follows, the above recipe will be referred to as the “standard procedure”. The weights, w_{ij} , only depend on the source grid and destination grid, so a single set of weights can be generated that can be applied in remapping any field. It should be noted that the approximate areas calculated by the remapping algorithm are of no interest once the mapping weights have been generated. In the subsequent mapping of a variable from the source grid to the destination grid, only the true areas of cells are needed.

360 As we have seen, the errors in remapping data to or from grids defined by longitude and latitude coordinates arise because the remapping algorithm constructs cell shapes with latitude sides that arc poleward following great circles rather than following latitude circles. The larger the longitudinal width of a cell, the larger its distortion. This suggests a simple strategy to improve the accuracy of the remapping algorithm. The destination grid and possibly the source grid (if its longitudinal resolution is coarse) should be subdivided into a number of sub-cell sectors, each of equal width. If the source cells have been subdivided,

365 then each source cell value is replicated across each of its sub-cells, as is each unmasked fraction. Next, the weights are obtained for mapping variables from the possibly subdivided source grid to the subdivided destination grid. Once the remapped field has been obtained for the fine-resolution destination grid, the sub-cells within each of the coarse-resolution destination cells are simply collected and averaged to obtain the final destination value. Each of the fine-resolution cells is contained within the boundaries of a single coarse destination cell, so no weighting is necessary.

370 Care must be taken when the standard procedure for remapping is applied to a variable representing conditions within layers of the atmosphere or ocean to ensure that mass-weighted means are preserved (as opposed to the usual area-weighted means). Additional complications might be encountered when a variable represents the ratio of two quantities (e.g., specific humidity



is the ratio of the mass of water vapor to the mass of air), where, rather than preserving the global mean ratio, it is better to preserve the two quantities themselves. The following guidelines may be helpful in treating these possibly troublesome variables.

(a) To remap a quantity representing a grid cell area fraction (e.g., cloud fraction, sea ice fraction, land fraction), the destination fractions should be calculated in the same way as unmasked fractions were calculated in step 4 above.

(b) To conserve the total energy flowing through the surface of the earth), the heat flux must be expressed as a flux per unit area (“flux density” with units of, for example, W/m^2 , not W). Then the standard procedure is followed to remap the flux density to the destination grid where it is scaled by c , as defined by (20).

(c) To remap the albedo (reflected radiation divided by incident radiation), which is undefined when the incident radiation is zero, it is best to conservatively remap the incident and reflected radiation flux densities (commonly termed “radiative fluxes”) and then form their quotient, rather than directly remapping the albedo. Destination values should be considered “missing” (undefined) where the remapped incident radiation is 0.

(d) There are applications where the total volume of a space (which might be partially masked) should be preserved. For grids constructed with height (or depth) used as a vertical coordinate, this can be achieved by calculating appropriate cell thicknesses on the destination grid. The standard procedure above is followed, applied to cell thickness. The resulting values must be scaled by c , as defined by (20). The unmasked portion of a destination cell volume is the product of the remapped cell depth, the destination cell fraction that is unmasked (f_{dj}), and the true destination cell area (A_{dj}^*).

(e) For most 3-d quantities, remapping should preserve the mass-weighted mean, rather than the area-weighted mean. Prior to remapping such variables, the mass field must be remapped conservatively. In order to preserve the total mass within a layer, the mass per unit area (M) of destination cells can be obtained following the standard remapping procedure. We consider two cases: models for which the bounds on layers can be expressed as a pressure and models for which the bounds on layers can be expressed as a distance.

When the layer pressure thicknesses can be determined, the mass per unit area in the layer is $M = \Delta p/g$ (where Δp is the pressure thickness of the layer, which may vary with longitude and latitude, and g is the acceleration due to gravity). Preserving the mass within a layer is equivalent to preserving the pressure thickness of the layer. This is achieved following the standard remapping procedure with $F_{si} = \Delta p_{si}$ and scaling the result by c , as defined by (20).

When the layer thicknesses can be determined, the layer mass per unit area is equal to the product of cell density (ρ) and layer thickness (Δz), so that standard procedure is applied to $\rho \Delta z_{si}$. Again, the result is scaled by c , as defined by (20). By this procedure we preserve global mass, but we should also like to define density and layer thickness on the destination grid such that their product is M_{dj} . For models with a uniform source-grid layer thickness, Δz , it makes sense for the source-grid thickness to carry over to the destination grid. Then the density is given by $\rho_{dj} = M_{dj}/\Delta z$. If, instead, density is uniform in a source grid layer, then $\Delta z_{dj} = M_{dj}/\rho$. If, however, both density and thickness vary on



- 405 the source grid, then one can choose whether to preserve the global mean layer thickness or the global mean density. One of these fields can be remapped, preserving its mean, and then the other calculated by dividing M_{dj} by the first field.
- (f) Once the mass per unit area is obtained for each destination grid cell, as just described in (e) above, the formula for preserving mass-weighted integrals or means can be derived. For example, to remap temperature (T) in a layer such that the total internal energy is conserved within a layer, remap the temperature, weighted by each grid cell's mass per unit area, and then divide by the cell mass per unit area on the destination grid (M_{dj}). For a pressure coordinate model with a layer thickness (Δp_{si}) that depends on location, the mass-weighted temperature $U = T\Delta p$, which is proportional to internal energy per unit area, is mapped to the destination grid following the standard procedure. The result, U_{dj} , is then divided by the pressure thickness on the destination grid (defined, as described in (e), such that the global mass is preserved), yielding the temperature field consistent with conservation of internal energy: $T_{dj} = U_{dj}/\Delta p_{dj}$. For a height or depth coordinate model, Δp is replaced by $\rho\Delta z$ in the above formulas, with care taken to preserve the global mass in the layer. Note that for layers of uniform mass thickness (either constant Δp or constant $\rho\Delta z$), there is no need to consider mass, and instead, the simpler procedure described in (b) can be applied directly without regard to layer thickness.
- 410
- 415
- (g) The amount of a substance in a layer of the atmosphere or ocean is often expressed as a ratio. To remap quantities of this kind, separately remap the quantities represented by the numerator and denominator and then form their ratio, as in the following examples:
- 420
- for specific humidity, q (mass of water vapor divided by mass of air containing the water), preserve separately the mass of water vapor and the mass of the air. First conservatively remap the water vapor mass per unit area ($q_{si}M_{si}$). Then remap the mass of air per unit area (M_{si}), as described in (e) above. Finally, form the ratio of the two remapped fields to obtain the specific humidity on the destination grid. A similar procedure can be applied in remapping any mass fraction.
 - for water vapor mixing ratio (mass of water vapor divided by mass of dry air), preserve separately the mass of the water vapor and the mass of dry air. A similar procedure can be applied in remapping any mass mixing ratio.
 - for number concentration (number of particles divided by volume), preserve separately the number of particles and the volume. Then form their ratio.
 - for mass concentration (mass of substance divided by total volume of mixture), preserve separately the mass of the substance and the volume. Then form their ratio.
 - for mole concentration (number of moles per unit volume of, for example, a chemical species in the atmosphere or ocean), preserve separately the moles of the substance and the volume. Then form their ratio.
 - for volume mixing ratio (number of moles of a constituent divided by number of moles of all constituents combined; sometimes referred to as mole fraction), preserve separately the number of moles of the constituent of interest and the number of moles of all constituents combined. Then form their ratio.
- 430
- 435



440 (h) In remapping relative humidity (mixing ratio divided by saturation mixing ratio), one might want to preserve the relationship between the remapped mixing ratio and the remapped temperature used to define saturation mixing ratio. That is, one might want the relative humidity on the destination grid to be defined by the ratio of a conservatively remapped mixing ratio divided by a saturation mixing ratio based on a conservatively remapped temperature.

5 Interpolating conservatively in the vertical

445 When remapping a 3-d field both vertically and horizontally, the vertical dimension must be handled carefully to preserve a global mass-weighted integral. We consider here the specific case of interpolating from a model's native vertical grid to surfaces of constant mass per unit area. If a field will ultimately be mapped to multiple horizontal destination grids, then performing the vertical interpolation first will usually reduce the computational expense. Once vertical interpolation is completed, the remapping to a different horizontal grid proceeds as described in the previous sections.

450 Compared with the generation of weights needed to remap conservatively in the horizontal, it is much easier to define the weights that will preserve integrals in the vertical. This is because the overlap of source and destination grid cells in the vertical is one dimensional, and only the cells thicknesses must be considered (not their shapes). For data stored on native model levels, bounds defining the vertical extent of each grid cell are an essential component of the grid definition and should be known. Furthermore, the pressures associated with those interfaces should be derivable. Then the mass per unit area contained within the upper and lower bound of a layer can be calculated by dividing the pressure difference across the layer by g . The weights can then be obtained by overlaying the pressure bounds from the native grid onto the destination grid bounds and 455 determining the fraction of each source cell that lies within the vertical extent of each destination cell. These fractions are used to remap the data in the vertical through matrix multiplication. A difference from horizontal remapping is that the weights are not uniform across the other dimensions of the data; they can vary from one location to another and may evolve over time (e.g., in the atmospheric surface layer where surface pressure may vary in time). It is therefore not possible to calculate the weights once and for all as is done in horizontal remapping. Fortunately, calculating the weights for interpolating in the vertical 460 is computationally much less demanding than in the horizontal, so remapping 3-d fields remains practical.

6 Summary and concluding remarks

465 Most conservative remapping packages (see Appendix A) generate mapping weights based on grid cells that may differ slightly in shape from the true cell shapes. Typically, a remapping algorithm will construct cell polygons with edges that follow great circles and then use these to determine cell areas and mapping weights. On the other hand, many models and analysis grids are constructed on spherical grids with grid cell bounds that follow lines of constant longitude and lines of constant latitude (not great circles). If data are mapped from or to a grid of this kind, the remapping algorithms can fail to preserve the true global mean or integral of a field. The algorithms preserve instead a global mean based on its approximate representation of cell shapes and areas, which generally differs from the true mean. This difference may especially matter when gauging whether



a model, having reached equilibrium, is conserving energy. If the global mean net top of the atmosphere (TOA) energy flux is
470 in fact zero, as evaluated based on the original grid and correct cell areas, remapping that data and calculating the mean on a
new grid could lead to a different conclusion.

Another limitation of many remapping packages is that although they may be able to treat gridded data where a binary mask
applies (e.g., screening regions of missing data or limiting analysis to the ocean or land regions alone), they are not designed to
conveniently handle data values that are representative of only a portion of a grid cell (i.e., are partially unmasked). Moreover,
475 often the easiest option offered for handling such cases is to perform the computationally intensive recalculation of weights
each time a new mask is imposed.

Here we have suggested how users can apply the weights provided by remapping packages, while avoiding or correcting
their limitations. For a given pair of source and destination grids, the remapping weights need only be calculated once; the
weights are independent of any full or partial masking of the source data. Each destination field can then be calculated via very
480 sparse-matrix multiplications. The recipes appearing in section 4 provide step-by-step instructions on how to handle various
cases. These recipes apply even when the remapping algorithm has correctly represented the shapes and areas of grid cells;
when that is true, the steps involving correction of the mean can be skipped.

Conservative remapping of the kind considered here must always operate on variables that are independent of the cell's
area. For example, rather than remap the area of snow cover in grid cells, the areas first must be converted to fractions, which
485 can be conservatively remapped and then converted back to areas. Most variables reported from models are intensive, so such
conversions are rarely necessary.

Conservative mapping is obviously required if it is important to preserve the global integrals (or means) of a field. When this
is not essential, other methods of interpolating data to a destination grid may lead to a more physically consistent and realistic
looking result. Consider, for example, the geopotential height and wind fields carried on a relatively coarse source grid. If these
490 fields were mapped conservatively to a much finer resolution grid, box-fill contour plots of the resultant fields would look like
slightly blurred versions of the box-fill plots of the original fields. The sub-cells wholly contained within a given source cell
would all share the same value; there would be no variation except for the relatively few cells at the borders of the original
source cells. Thus, within the confines of each original source cell, the geopotential height and winds would be constant, and
at the borders of original source cells, there would be large gradients. With non-zero wind values but no geopotential gradients
495 within the confines of the original cell, the geostrophic balance generally prevailing outside the tropics would be upset. In
general, when mapping from a coarse to a fine grid, simple bilinear interpolation should lead to more realistic gradients and
more realistic looking results than conservative remapping.

We have shown that with available remapping packages, careful application of the remapping weights, unmasked fractions,
and in some cases layer masses can result in reasonably accurate results that preserve the true global integrals or means of
500 interest. If grid cell shapes were correctly recreated by the remapping algorithms, the remaining inaccuracies could be reduced
further and the standard remapping procedure could be simplified (e.g., by eliminating the corrections needed to adjust the
global means). This would seem to provide strong motivation for augmenting remapping packages with the option to correctly
construct the commonly encountered longitude by latitude grids with cell edges conforming to the true grid cell shapes.



505 *Code availability.* The calculations performed as part of this research were based on simple applications of the equations appearing in the manuscript, which were applied to completely artificial data (six temperatures that were assumed to decrease linearly from equator to pole).

Appendix A: Remapping packages

The focus here has been on the second step of conservative remapping procedures: applying remapping weights that have been generated by the first step. In climate modeling, among the most commonly used packages used to generate these weights are the following:

- 510 – C-Coupler2: This package was developed for use in coupling components of climate models (Liu et al., 2018) and includes a conservative remapping capability.
- Climate Data Operators (CDO): This package, designed to manipulate and analyze climate and weather prediction data, includes remapping capabilities based on the YAC package (see below). Documentation is available at <https://code.mpimet.mpg.de/projects/cdo/wiki>.
- 515 – Earth System Modeling Framework (ESMF) Regrid Weight Generator (ERWG): This library contains a number of packages useful in the analysis of climate data, including a conservative remapping capability. Documentation is available at <https://earthsystemmodeling.org/regrid/>.
- Model Coupling Toolkit (MCT): This package includes an application of the Jones (1999) method that has been parallelized (Larson et al., 2005). It is used by the OASIS coupler, and by a number of climate models.
- 520 – netCDF Operators (NCO): This toolkit manipulates and analyzes data of interest to the geophysical community and includes two options for creating remapping weights: TempestRemap and ESMF (ERWG). Documentation is available at <https://sourceforge.net/projects/nco/>, <https://nco.sourceforge.net/nco.pdf>, and <https://acme-climate.atlassian.net/wiki/spaces/DOC/pages/754286611/Regridding+E3SM+Data+with+ncremap>
- 525 – OASIS: This software was developed for coupling components of climate models and in its latest version, it relies on MCT to generate remapping weights (Craig et al., 2017). Documentation is available at http://www.cerfacs.fr/oa4web/oasis3-mct_4.0/oasis3mct_UserGuide/node42.html and http://www.cerfacs.fr/oa4web/oasis3-mct_4.0/oasis3mct_UserGuide/node48.html#subsec_mapdata.
- Spherical Coordinate Remapping and Interpolation Package (SCRIP): This is the first library to implement the Jones (1999) approach to remap conservatively (see <https://github.com/SCRIP-Project>). Most other conservative remapping
- 530 packages have been based on SCRIP.
- TempestRemap: This is a conservative, consistent and monotone remapping package for arbitrary grid geometry with support for finite volumes and finite elements (see Ullrich and Taylor, 2015; Ullrich et al., 2016).



- Yet Another Coupler (YAC): A conservative remapping algorithm is included in this climate model component coupler (Hanke et al., 2016).

535 *Author contributions.* The author was solely responsible for this contribution.

Competing interests. The author declares that he has no competing interests.

Acknowledgements. I thank Paul Durack and Paul Ullrich for their helpful comments. This work was performed under the auspices of the U.S. DOE by Lawrence Livermore National Laboratory under contract DEAC52-07NA27344. It was supported by the Regional and Global Modeling Analysis (RGMA) program area under DOE's Biological and Environmental Research (BER) Program.



540 References

- Craig, A., Valcke, S., and Coquart, L.: Development and performance of a new version of the OASIS coupler, OASIS3-MCT_3.0, Geoscientific Model Development, 10, 3297–3308, <https://doi.org/10.5194/gmd-10-3297-2017>, 2017.
- Hanke, M., Redler, R., Holfeld, T., and Yastremsky, M.: YAC 1.2.0: new aspects for coupling software in Earth system modelling, Geoscientific Model Development, 9, 2755–2769, <https://doi.org/10.5194/gmd-9-2755-2016>, 2016.
- 545 Jones, P. W.: First- and second-order conservative remapping schemes for grids in spherical coordinates, Monthly Weather Review, 127, 2204–2210, [https://doi.org/10.1175/1520-0493\(1999\)127<2204:FASOCR>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<2204:FASOCR>2.0.CO;2), 1999.
- Larson, J., Jacob, R., and Ong, E.: The model coupling toolkit: A new Fortran90 toolkit for building multiphysics parallel coupled models, The International Journal of High Performance Computing Applications, 19, 277–292, <https://doi.org/10.1177/1094342005056115>, 2005.
- Liu, L., Zhang, C., Li, R., Wang, B., and Yang, G.: C-Coupler2: a flexible and user-friendly community coupler for model coupling and
550 nesting, Geoscientific Model Development, 11, 3557–3586, <https://doi.org/10.5194/gmd-11-3557-2018>, 2018.
- Ullrich, P. A. and Taylor, M. A.: Arbitrary-order conservative and consistent remapping and a theory of linear maps: Part I, Monthly Weather Review, 143, 2419–2440, <https://doi.org/10.1175/MWR-D-14-00343.1>, 2015.
- Ullrich, P. A., Devendran, D., and Johansen, H.: Arbitrary-order conservative and consistent remapping and a theory of linear maps: Part II, Monthly Weather Review, 144, 1529–1549, <https://doi.org/10.1175/MWR-D-15-0301.1>, 2016.




Figure 1. For the example described in the text, destination grid cell values resulting from different remapping options. The source data were defined on a longitude by latitude grid of 30° by 15° resolution and then mapped to a destination grid with half the longitudinal resolution but the same latitude spacing (i.e., 60° by 15°). The solid black line (source grid ‘truth’) also represents the destination values that would result from correctly remapping the data to the destination grid, applying an algorithm that correctly reconstructed the grid cell shapes.

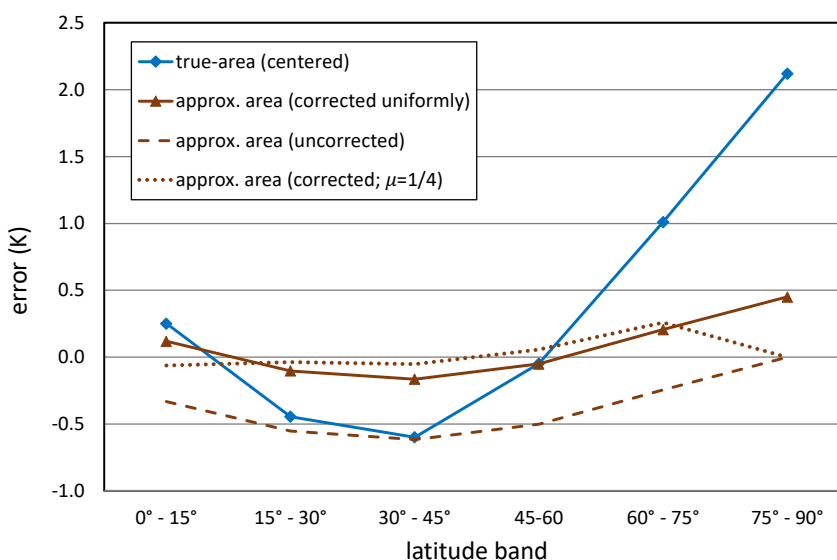


Figure 2. For the example described in the text, error in destination grid cell values resulting from different calculational options.




Figure 3. For the example described in the text, dependence on the longitude cell width of the RMS error in destination cell values, with the error calculated over all latitudes and weighted by area. The RMS errors have been normalized by the true spatial standard deviation of the variable. Expressed in this way, an error equal to 1 means, for example, that the RMS error is as large as the spatial standard deviation of the variable, which in this example is 7.2 K. The mapping is always from a source grid with longitude cell widths half that of the destination grid but the same latitude resolution (15° latitude bands for all longitudinal resolutions).




Figure 4. For the example described in the text, dependence of the temperature correction, $(\gamma^\mu / \bar{\gamma}^\mu)(\bar{T}_d - \bar{T}_s)$, on temperature and the exponent μ . Curves shaped the same apply to any variable with the abscissa spanning the full range of the source data and the ordinate appropriately rescaled.



Table 1. True cell areas (A^*) and approximate cell areas (A) for source and destination spherical coordinate grids with longitudinal cell widths of 30° and 60° , respectively. The approximate areas are calculated assuming all bounds coincide with great circles. All areas are expressed as solid angles obtained by dividing the actual cell area by the square of Earth’s radius. For ease of comparison with destination cell areas, source cell areas have been doubled.

latitude band	source cells		destination cells	
	$2A_{si}^*$	$2A_{si}$	A_{dj}^*	A_{dj}
$0^\circ - 15^\circ$	0.271	0.277	0.271	0.297
$15^\circ - 30^\circ$	0.253	0.256	0.253	0.265
$30^\circ - 45^\circ$	0.217	0.216	0.217	0.213
$45^\circ - 60^\circ$	0.166	0.163	0.166	0.152
$60^\circ - 75^\circ$	0.105	0.101	0.105	0.090
$75^\circ - 90^\circ$	0.036	0.034	0.036	0.030