


# Converting the OMWG diagnostic scripts to NCL

Observations and lessons learned

Dave Brown

NCL Development Team

# Topics



- Motivation for this project
- Background and current status
- Implementation guidelines
- Comparison of graphical output
- NCL vs. IDL
- Comparison of source code
- Lessons learned

# Motivation


- ParVis: provide parallel-processing solutions for the big data problem facing climate researchers
- Led by ANL; collaborators include NCAR, PNNL, SNL and UC-Davis
- Multiple goals including (among others):
  - (long term) ParNCL: a parallel version of NCL
  - (short term) Use SWIFT, a task-parallel scripting tool to improve performance for existing tasks
- The diagnostics make good ParVis case studies
- **Immediate benefit: provide non-proprietary open and free code that users can deploy anywhere**

# AMWG diagnostics status

- C-shell scripts run NCO tools for data reduction and NCL for analysis and viz
- Converted to Swift originally by John Dennis
- Changes to Swift to accommodate the diagnostic package work flow.
- For ParVis, an all-NCL version developed for comparison between the new ParNCL and the existing version using NCO tools (or a parallel-enabled replacement, Pagoda from PNNL)

# Various

## Datasets



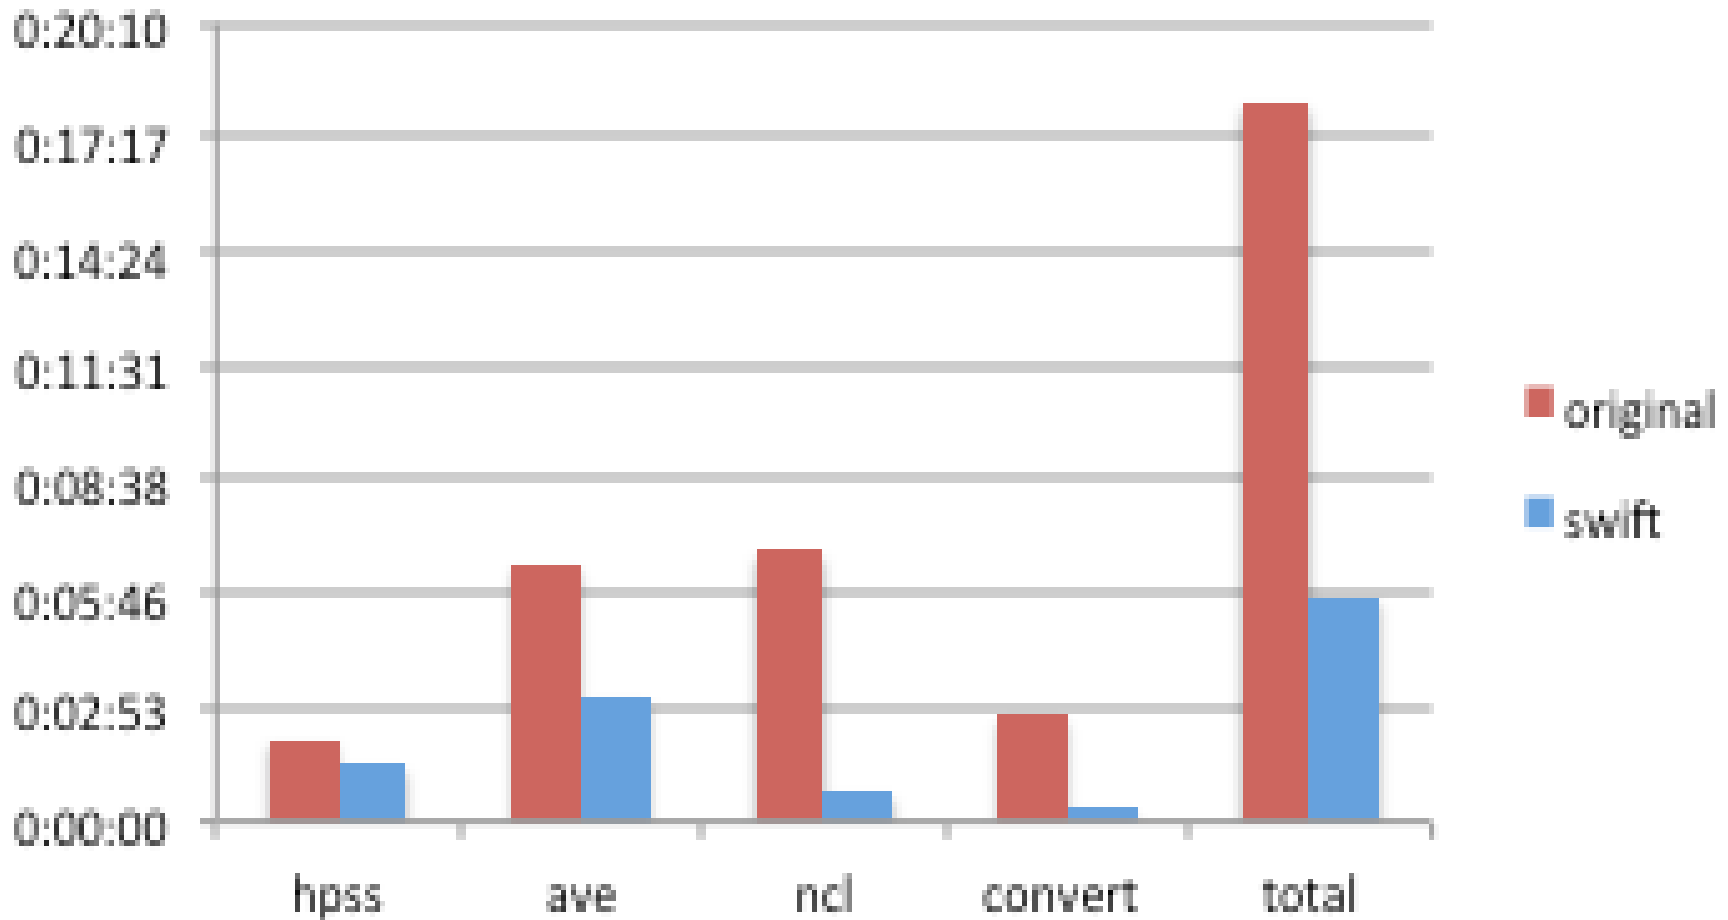
Courtesy Mike Wilde

## AMWG Diagnostics

# OMWG diagnostics status

- 96% complete – 84 of 87 scripts
  - popdiag and popdiagdiff finished
  - popdiagts: 3 to go
- 2 of 3 Fortran procedures (wrapped as shared objects for now – eventually will become built-in NCL routines)
- Basically transparent to user – scripts work the same as they always have

# Preliminary Timings for popdiag.csh



Courtesy Sheri Mickelson

# Conversion project guidelines



- Conservative approach
- Results must match mathematically and graphically
- Therefore initial version retains original colormaps, contour levels, and line colors for ease of verification
- Similar positioning of annotations, but font styles, etc. allowed some variation
- Fairly literal translation of code where performance not affected
- Array arithmetic used more aggressively since NCL looping performance is slower





# OMWG diagnostic output comparisons

- NCL output online:
  - [http://www.ncl.ucar.edu/Applications/popdiag/pd.1981\\_2005/popdiag.html](http://www.ncl.ucar.edu/Applications/popdiag/pd.1981_2005/popdiag.html)
- Current IDL output online:
  - [http://www.cesm.ucar.edu/experiments/cesm1.0/diagnostics/b40.20th.track1.1deg.005/ocn\\_1981-2005-obs/popdiag.html](http://www.cesm.ucar.edu/experiments/cesm1.0/diagnostics/b40.20th.track1.1deg.005/ocn_1981-2005-obs/popdiag.html)


IDL:




NCL:



IDL:





( 3.06e+01 to 3.61e+01 by 0.40psu)  
(MODEL - OBS)


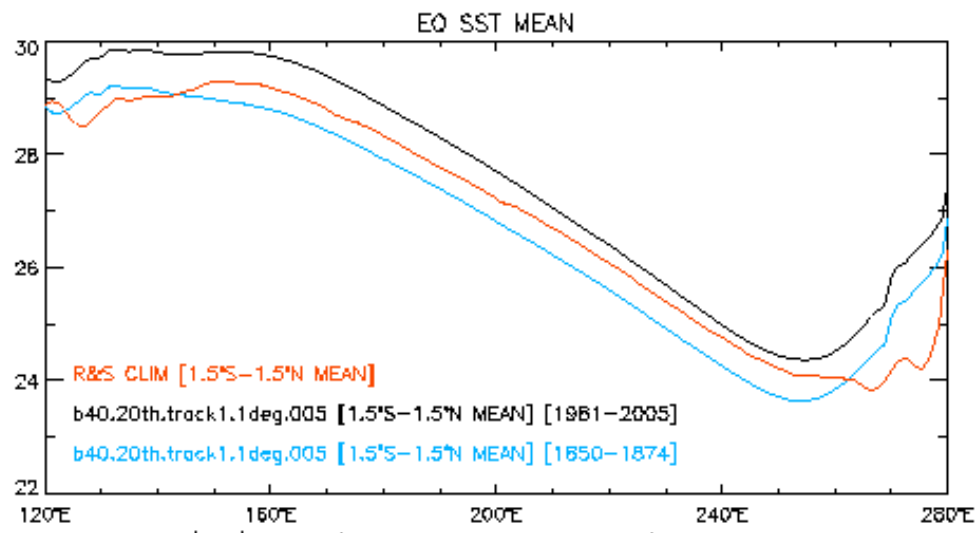


(-1.63e+00 to 1.01e+00 by 0.20psu)

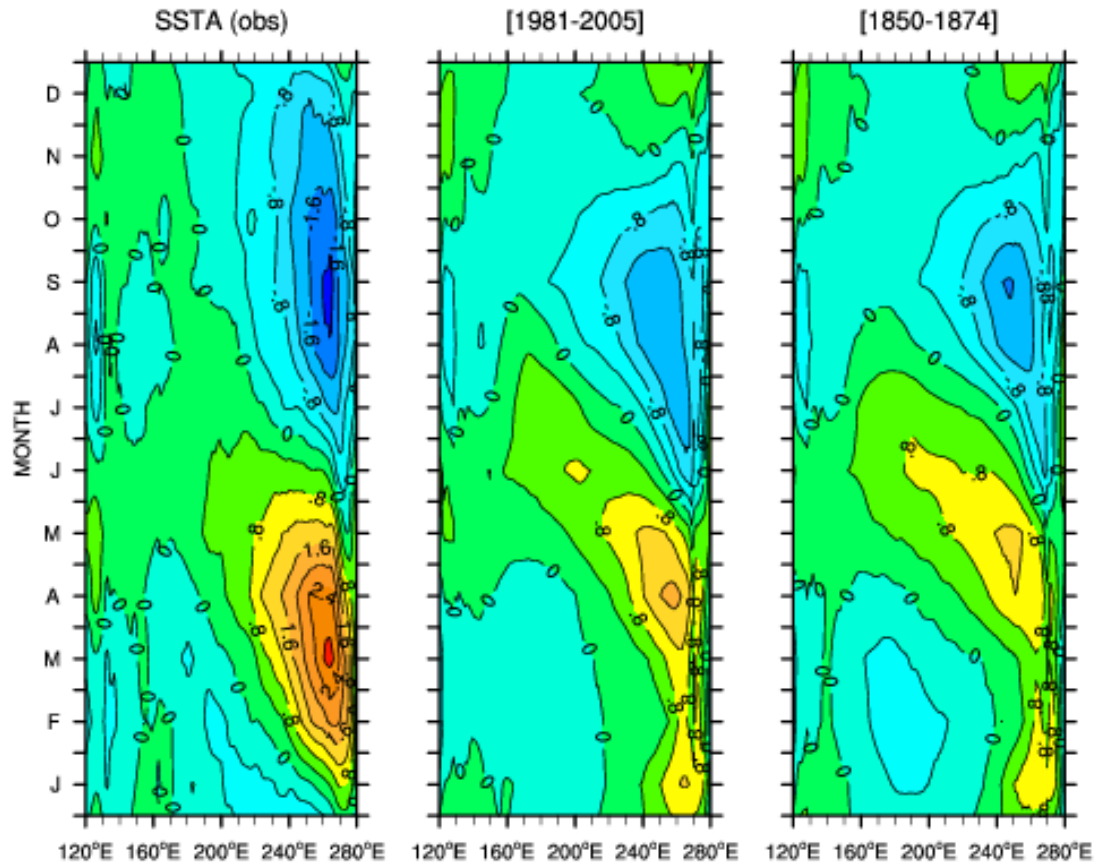

NCL:



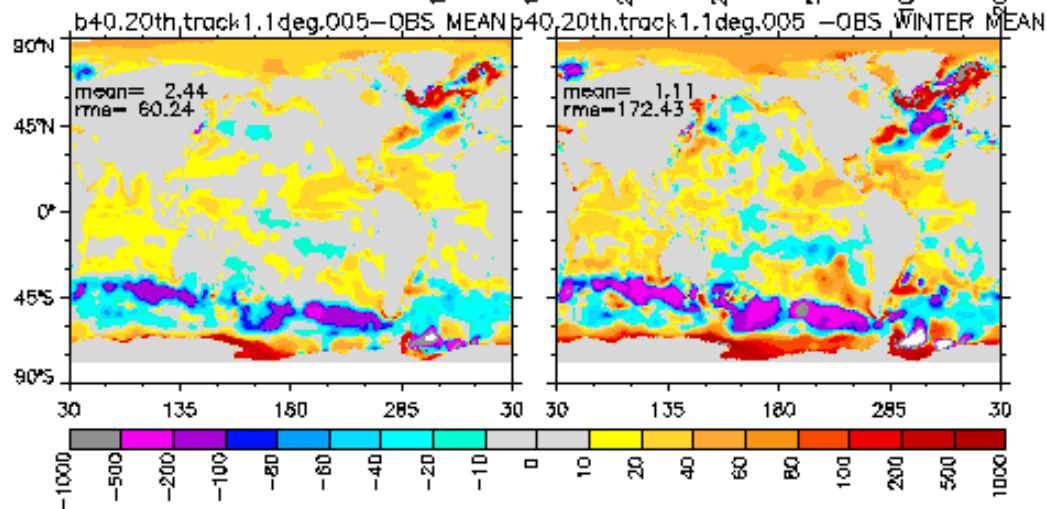
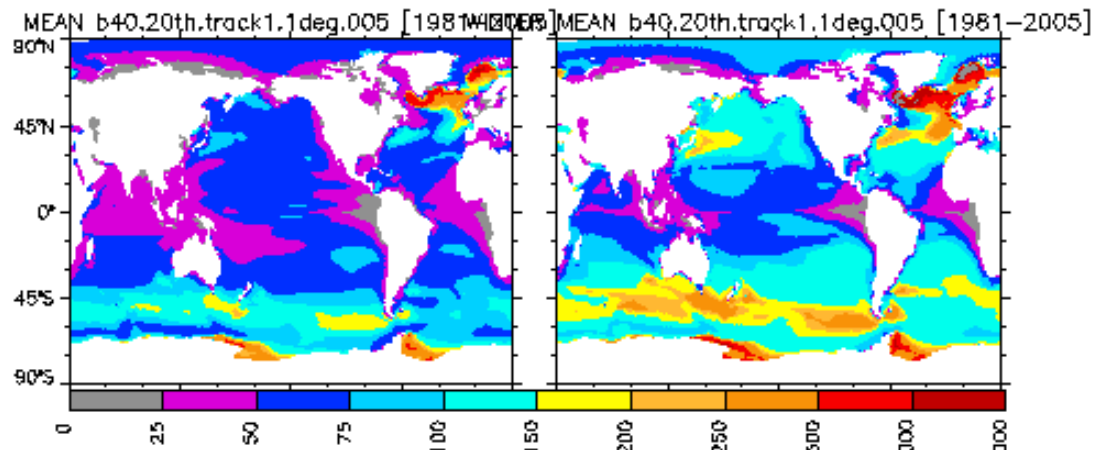
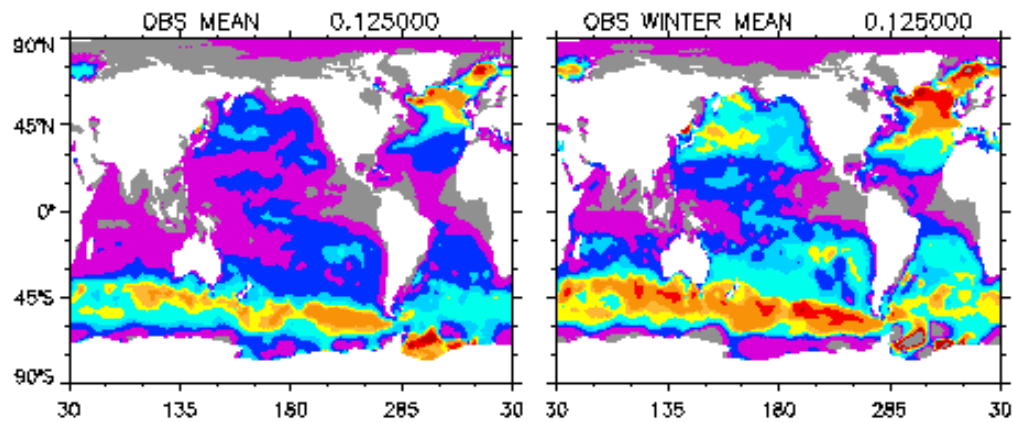
IDL:



NCL:

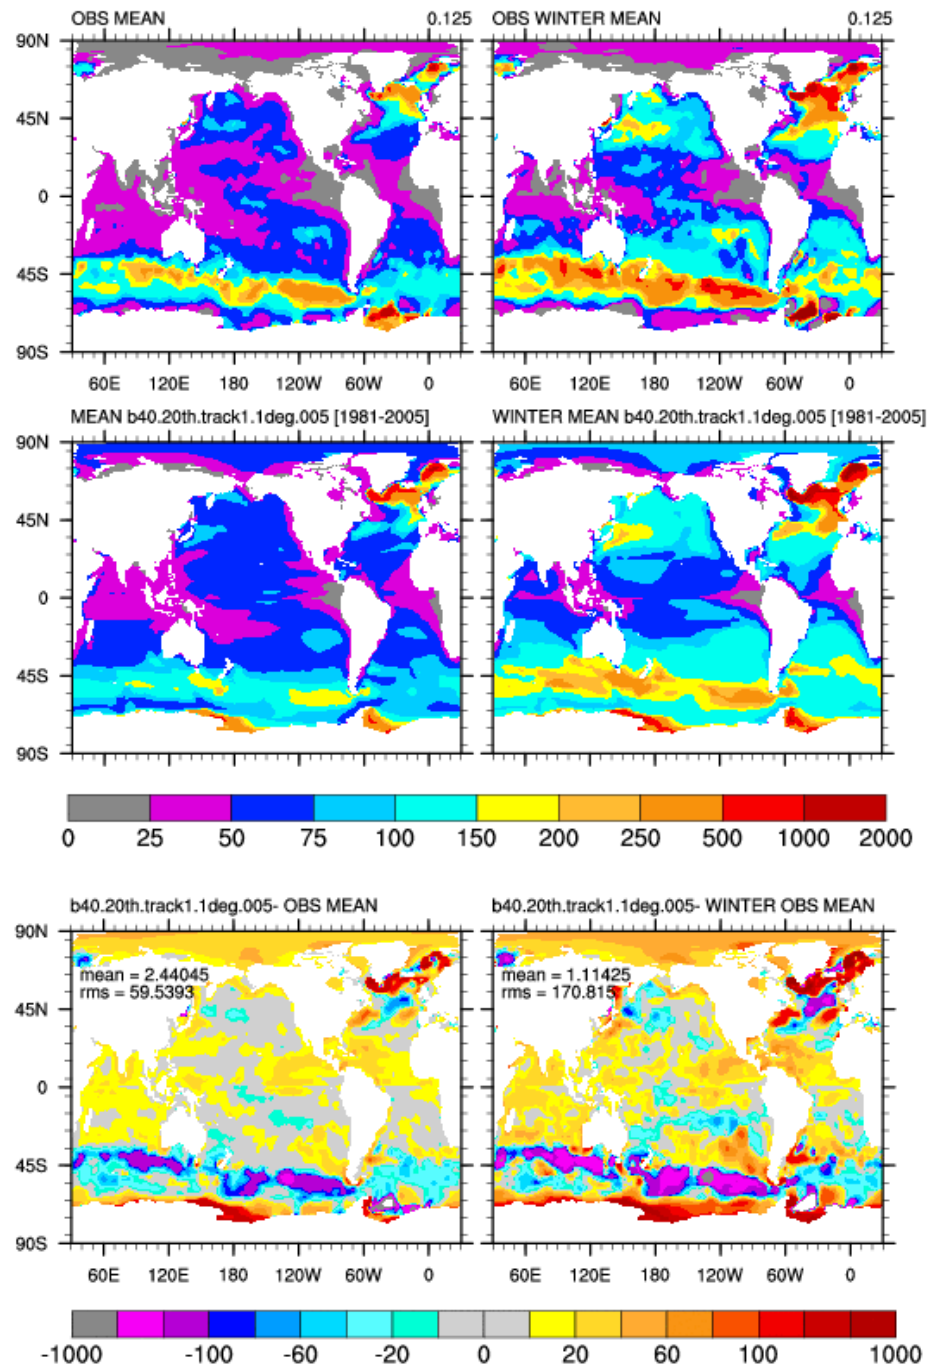


IDL:





NCL:



# NCL vs. IDL

## ( the good, the bad, and the ugly)

- Many apparent similarities
  - ; (semicolon) starts a comment
  - Fortran-like syntax features: e.g. .eq. (NCL), eq (IDL)
  - Overall verbosity (lines of code): 14424 (NCL), 14388 (IDL)
  - Similar array syntax: 0-based element counting
- Significant differences
  - NCL: row-major like C; IDL: column-major like Fortran
  - Graphics code has a different model
  - NCL's built-in support for missing values helps simplify code
  - NCL's NetCDF-like variable model allows easier access for attributes and other metadata
  - IDL looping is definitely faster (script is compiled)
  - (Therefore) more important to use array syntax in NCL

# NetCDF file handling comparison

```
; IDL open file, read variable, and handle attribute if it exists
fileid = ncdf_open(file_netcdf)
varid = ncdf_varid(fileid, 'SALT')
ncdf_varget, fileid, varid, salt
f_struct = ncdf_varinq(fileid,varid)
n_att = f_struct.natts
for n_att=0,n_att-1 do begin
  if ( ncdf_attname(fileid, varid, n_att) eq 'scale_factor' ) then begin
    ncdf_attget, fileid, varid, 'scale_factor', scale_field
    good = where(salt gt -10. AND salt lt 1.e10)
    salt[good] = scale_field * salt[good]
  endif
endifor
```

---

---

```
;NCL open file, read variable and handle attribute if it exists
;Note: attribute is part of variable, _FillValue support ensures that
;missing values are automatically ignored
```

```
fileid = addfile(file_netcdf,"r")
salt = fileid->SALT
if (isatt(salt,"scale_factor")) then
  salt = salt * salt@scale_factor
end if
```

(the good)

# Calculating weighted average (IDL)

```
; variable field contains temperature anomalies : lon x lat x time  
; the task is to average values near the equator from y_min to y_max  
; tarea has the area weights on the T grid  
; anom is lon x time averaged over lat  
; triple-nested loop handles each array element individually
```

```
anom = dblarr(nx,nt)  
anom(*,*) = double(0.)  
for n=0,nt-1 do begin  
  for i=0,nx-1 do begin  
    area_wt = double(0.)  
    max_anom = double(0.)  
    for j=y_min,y_max do begin  
      if ( field(i,j,n) lt missing ) then begin  
        anom(i,n) = anom(i,n) + tarea(i,j) * field(i,j,n)  
        if (anom(i,n) gt max_anom) then max_anom = anom(i,n)  
        area_wt = area_wt + tarea(i,j)  
      endif  
    endfor  
    if ( area_wt ne 0. ) then begin  
      anom(i,n) = anom(i,n) / area_wt  
    endif else begin  
      anom(i,n) = missing  
    endelse  
  endfor  
endfor
```

## Calculating weighted average (NCL)

```
; variable field contains temperature anomalies : time x lat x lon
; the task is to average values near the equator from y_min to y_max
; tarea has the area weights on the T grid
; anom is time x lon averaged over lat
; conforming the dimensions of tarea with the field variable allows
; NCL to perform element by element array arithmetic and avoids loops
; However, note that the conform_dims function creates an array with nt
; redundant copies of the same data. The temporary array then needs to be
; deleted.
```

```
sub_y = y_max - y_min + 1
tarea_conform = conform_dims((/ nt, sub_y, nx /), \
                             tarea(y_min:y_max,:), (/ 1, 2 /))
subfield = tarea_conform * field(:,y_min:y_max,:) ; time * lat * lon
anom = dim_sum_n_Wrap(subfield,1)
tarea_anom = dim_sum_n_Wrap(tarea_conform,1)
anom = anom / tarea_anom
delete(tarea_conform)
delete(subfield)
```

(the bad)

# Smoothing code for mixed layer depth value (IDL)

```
; a more complicated code with multiple nested loops that requires access to  
; adjacent cells along 2 dimensions during each pass.  
; Only the beginning shown here
```

```
for ns=1,ns_max do begin
```

```
  print, ' smoothing pass .... ', ns
```

```
  field_temp_1 = MLD
```

```
  for j=1,ny-2 do begin
```

```
    for i=0,nx-1 do begin
```

```
      im1 = i-1
```

```
      ip1 = i+1
```

```
      if ( i eq 0      ) then im1 = nx-1
```

```
      if ( i eq nx-1 ) then ip1 = 0
```

```
        cc = double(tarea(i  ,j  ))
```

```
        ce = double(tarea(ip1,j  ))
```

```
        cw = double(tarea(im1,j  ))
```

```
        cn = double(tarea(i  ,j+1))
```

```
        cs = double(tarea(i  ,j-1))
```

```
        sum = cc + ce + cw + cn + cs
```

```
        cc = cc / sum
```

```
        ce = ce / sum
```

```
...
```

# Smoothing code for mixed layer depth value (NCL)

; Sample lines of my attempt to recreate this code in NCL eliminating loops.  
; Eventually I got it to work more or less, but it still did not have the  
; desired performance and it just looks too complicated to be maintainable.

```
tarea_sum = tarea
tarea_sum(1:ny-2,1:nx-2) = \
    tarea(1:ny-2,1:nx-2) + tarea(1:ny-2,:nx-3) + tarea(1:ny-2,2:nx-1) + \
    tarea(:ny-3,1:nx-2) + tarea(2:ny-1,1:nx-2)
tarea_sum(1:ny-2,0) = \
    tarea(1:ny-2,0) +tarea(1:ny-2,nx-1) + tarea(1:ny-2,1) + \
    tarea(:ny-3,0) + tarea(2:ny-1,0)
...

MLD_new(:, :, 1:ny-2, 1:nx-2) = \
    MLD(:, :, 1:ny-2, 1:nx-2) * cc_c(:, :, 1:ny-2, 1:nx-2) + \
    MLD(:, :, 1:ny-2, :nx-3) * cw_c(:, :, 1:ny-2, 1:nx-2) + \
    MLD(:, :, 1:ny-2, 2:nx-1) * ce_c(:, :, 1:ny-2, 1:nx-2) + \
    MLD(:, :, :ny-3, 1:nx-2) * cn_c(:, :, 1:ny-2, 1:nx-2) + \
    MLD(:, :, 2:ny-1, 1:nx-2) * cs_c(:, :, 1:ny-2, 1:nx-2)      (the ugly)
; etc.
```

When the effort to avoid looping in NCL means the code starts looking like this, it's probably time to switch to Fortran and create a shared object.

# Summary

- New OMWG diagnostic suite verified and available by the end of the year
- Freely distributable open source
- Performance and graphics similar to existing suite
- Future improvements possible
- Suggestions welcome



Thanks to Susan Bates and Gokhan Danabasoglu for their support of this project and Dennis Shea and Mary Haley for their good advice

Questions?

Email: [dbrown@ucar.edu](mailto:dbrown@ucar.edu)

