New Options for POP2 Time-Averaged History Files

Nancy Norton
NCAR
December 2009

Background

- LANL POP2 time-average output file features:
  - 2D and 3D fields only: i,j(,k) indices
  - All fields of the same frequency
  - No time dimension; no time series
- Additional CCSM POP2 TAVG output-file features:
  - Unlimited time dimension (netCDF)
  - C-F Compliance
  - Coordinate variables
  - Time-invariant fields
  - Nonstandard dimensions; 4- and 5-dimensional fields
  - Timestamp in filename (yyyy-mm-dd,yyyy-mm,yyyy)

CCSM TAVG Output-File Limitations

- Each file contained all selected fields written at a single frequency
- Each file contained many time-invariant fields (potentially large overhead)
- Very large files created by ecosystem model or at high resolution limited the number of output fields that could be selected

More Flexibility and Expanded Functionality December 2009

- Multiple TAVG output file “streams” (maximum nine; CCSM default configuration is three)
- One-time header option (write time-invariant info only once, at the beginning of each run segment)
- New frequency option, “once”
- Allow offset date for averaging interval (eg, weekly files yyyy-mm-07, yyyy-mm-14,…)
- New default filename conventions and new filename option for streams:
  - *.h.* (no change; base model, standard monthly output)
  - *.h.freq_string.* (base model, other frequencies, eg *.h.nday5’
  - *.h.tracer_string.freq_string.* (extra-tracer modules, eg*.h.ecosys.nmonth1)
- Timeseries !

More Flexibility and Expanded Functionality December 2009

- Controls for all *fields* within each TAVG output “stream” are customizable, and include:
  - Frequency
  - Start
  - Offset-date
  - One-time header
- Controls for each TAVG output “stream” *file* include:
  - Frequency
  - Name
  - Format (binary or netCDF)
  - TAVG restart file
- Extra-tracer modules can create their own streams or use the base-model streams

CCSM TAVG Default Setup

- Three output streams:
  - Monthly, daily, and one-time field frequencies
  - All files include full header information
  - Daily-frequency fields are bundled into monthly files
  - No offset dates
  - netCDF file format

TAVG Contents File

- New format: (i1,2x,a)
- First character indicates stream number or is a comment
- Duplicate entries ignored
- Unique fields counted; no need to specify the number of fields in the contents file
- Example:

```
#  An integer [1,9] in the first column identifies the field's output
#  stream number.
#  To deactivate a field, replace the output stream number with  # or !
1  UVEL
1  VVEL
1  WVEL
1  HMXL
2  HMXL_2
1  XMXL
2  XMXL_2
```

TAVG Namelist

- New fields
- Array entries
- Define up to nine array elements

Original TAVG Namelist

```
&tavg_nml
   tavg_freq_opt             = 'nmonth'
   tavg_freq                 = 1
   tavg_start_opt            = 'nstep'
   tavg_start                = 0
   tavg_fmt_in               = 'nc'
   tavg_fmt_out              = 'nc'
   tavg_contents              = 'tavg_contents'
   tavg_infile               ='./CASE.pop.hrestart.end'
   tavg_outfile               ='.CASE.pop.h'
  ltavg_nino_diags_requested  = .true.
```

New TAVG Namelist

```
&tavg_nml
   n_tavg_streams                 = 4
   ltavg_streams_index_present = .true.
   tavg_freq_opt                   = 'nmonth' 'nday' 'once' 'nday'
   tavg_freq                       = 1 1 1 5
   tavg_file_freq_opt              = 'nmonth' 'nmonth' 'once' ‘nyear’
   tavg_file_freq                  = 1 1 1 1
   tavg_stream_filestrings         = ' ' 'nday1' 'once' 'ecosys.nday5'
   tavg_start_opt                  = 'nstep' 'nstep' 'nstep' 'nstep'
   tavg_start                      = 0 0 0 0
   tavg_fmt_in                     = 'nc' 'nc' 'nc' 'nc'
   tavg_fmt_out                    = 'nc' 'nc' 'nc' 'nc'
   tavg_contents                   = 'tavg_contents'
   ltavg_nino_diags_requested  = .true.
   tavg_infile                    ='./CASE.pop.hrestart.end'
   tavg_outfile                   ='.CASE.pop.h'
   ltavg_has_offset_date          = .false. .false. .false. .true.
   tavg_offset_years              = 1 1 1 1
   tavg_offset_months             = 1 1 1 1
   tavg_offset_days               = 2 2 2 1
   ltavg_one_time_header          = .false. .false. .true. .false.
```

Feedback 2008/Status 2009

- ✓ Maintain backwards compatibility
- ✓ Who needs netCDF, anyway? (but note some features, like timeseries, are not supported with binary I/O)
- ✓ Who needs time-invariant information in files? Make inclusion of time-invariant grid info optional
- ✓ Parallel I/O developments may affect our decisions
- ✓ Memory may limit the number of fields a user can select at one time (multiple frequencies and multiple options per field)
- Regional diagnostics?
  - Regional diagnostics can be constructed during post-processing
  - Regional diagnostics might include high-frequency, upper-ocean sampling in the Tropics
- Can LANL HYPOP vertical remapping code be used?
- Counterproposal to “streams” that uses a single, xml-type file
- Allow run-time selection of tavg method (min,max,avg; presently hard-coded)
- Allow run-time selection of vertical levels stored

Summary 2009

- ✓ Significant progress towards increased TAVG flexibility/functionality goals
- ✓ Streams are being used in high-res runs
- ✓ Three streams are now default for CCSM base model: monthly, daily timeseries, and one-time
- ✓ Additional, independent streams can be used by extra-tracer modules
- ✓ Pio library implemented “under the covers” in io_netcdf.F90
- ✓ Additional streams error-checking needed