A Guide to the Workstation Eta

Matthew E. Pyle
Nonhydrostatic version
Revised: March 2001
mpyle@ncep.noaa.gov

Initial setup

1) Extract the contents of the tar file with the following command:

tar xf worketanh.tar

This will extract the contents of this file into a directory named *worketa_nh*.

2) Go to the *worketa_nh* directory and set up a few items that are machine dependent by running the *setup* script with either DEC, LINUX, or NORMAL as an argument for the script. NORMAL should cover all users not using DEC or LINUX machines.

EXAMPLE: If setting up the model on an HP workstation, then the command is **setup NORMAL**.

3) Create the software libraries used when compiling other codes. These compilations *require* a FORTRAN 90 compiler.

a) go to *worketa_nh/libraries/bacio.source* (note: portions of pathnames before *worketa_nh* will not be included in this manual, as they are unique to each user).

• define the computer system being used in the file *clib.h*.

For machines not included in the list (e.g., Sun) try the SGI and HP settings. The compilations in steps 5 and 11 link with the *bacio* library and are the real test of whether an appropriate version has been created.

• Edit the file *comp.com* to uncomment (lines beginning with # are treated as comments by the script) one of the three options for compiling the C portion of the code. Special compiler flags are required to compile on HP and DEC machines. Once the file has been edited, enter **comp.com** to create the file *bacio* in *worketa_nh/libraries*.

b) go to *worketa_nh/libraries/* and enter the command **make_ip_w3** to create the *w3lib* and *iplib* libraries.

c) go to the *dummyMPI* subdirectory.

The makefile used to build this library depends on the workstation type being used. For HP workstations, enter **make -f Makefile.hp**; for IBM workstations, enter **make -f Makefile.ibm**; for LINUX machines use **make -f Makefile.linux**; others can simply enter **make**.

4) Edit the file *make.inc* in *worketa_nh/dprep/src/configure* to specify the compiler name (FC) and compiler flags (FFLAGS, often used to set the level of optimization). Examples for certain machines are contained in this directory.

5) Change directories to *worketa_nh/dprep/install* and enter the command: **build_dprep**

This script will configure the paths in the makefile and then build the dprep (data preparation) executables.

6) Extract the topography data

If the topography data have been downloaded, they should now be placed into *worketa_nh/eta/static/topo*. If not, the data can be retrieved from *http://sgi62.wwb.noaa.gov:8080/wrkstn_eta/*. The topography data are tarred into 10° thick latitudinal strips extending from 180°W to 40°W, with separate strips providing coverage between 10°N and 80°N. Users modeling outside of this region need to request data via e-mail (e-mail: *mpyle@ncep.noaa.gov*).

The downloaded files should be uncompressed and untarred into individual topography files. The file names (e.g., *U20N130W*) give the coordinates of the southwest corner of the 10° × 10° region covered by the file.

7) The package defaults to use *US_2m_slm.ieee*, a 2' × 2' land/sea mask dataset covering from 10°N, 180°W to 80°N, 30°W. The other two datasets (*global_4m.ieee, global_8m.ieee*) are global and have 4 minute and 8 minute resolution, respectively. If using the 4 or 8 minute data, uncompress the appropriate file in *worketa_nh/eta/bin* and remember to specify it when editing the seares value in the *ETAIN* file (step #1 of "Setting up and running the model") below.

8) Change directories to *worketa_nh/eta/src/configure* and edit the *make.inc* file to set the compiler names and options. Examples from several workstation platforms are included in this directory. The -DLITTLE flag is critical for DEC and LINUX users, and the -DDEC flag is needed for DEC users.

9) Change directories to *worketa_nh/eta/install* and edit the *set_parmeta* script. This script sets information later used to compile the model, and installs some resolution-dependent files.

IM: Number of mass grid points along the first row, essentially half the total number of grid points in the west-east direction due to the horizontal staggering of mass and wind points (i.e., an IM=50, JM=99 grid gives a square domain).

JM: Number of rows in the north-south direction (NOTE: MUST BE ODD).

LM: Number of vertical levels. If running with the Eta vertical coordinate it must be set to 38, 45, 50, or 60 to work with the deta and radiation files included with the package. If running with the sigma vertical coordinate it must be set to 25, 30, 38, or 45 to work with the included radiation files.

LSM: Controls the number of output levels for isobaric data (discussed below in "The post-processor" section).

PT: The top pressure of the model atmosphere; 25 hPa is the value used at NCEP, and it must be set to 25 or 50 to work with the included radiation files.

HVAL: Character string used for symbolic links between files and FORTRAN unit numbers. Most machines need HVAL="fort.", but HP uses HVAL="ftn".

INPES and JNPES define how the model grid will be distributed over multiple processors (the model integrates on INPES $\times$ JNPES processors; one or more additional processors may dedicated to collecting and writing the output files). Users running the model on a 1-CPU machine or a machine without MPI should set these equal to one.

10) Go to *worketa_nh/eta/install* and enter the command: **buildall**

This script will set paths, use the information from the *set_parmeta* file to configure several items, and ultimately compile the model code, codes used to initialize the model, and codes that post-process model output. The optimizations set in step 8 are used for all of these steps.

11) You are now done with the initial set up. Changing the model dimensions (IM, JM, LM) or the number of CPUs (INPES, JNPES) requires repeating steps 9 and 10.

Getting the input data

1) Go to *worketa_nh/dprep/bin* and decide on an input data set. The workstation Eta can be initialized from global AVN data (full AVN or reduced-resolution WAFS form of the AVN), and from Eta data on grids 104, 212, and 221 (in tiled form). The area covered by the Eta output grids and the operational 22 km Eta domain is shown in *regions.gif*. Scripts have been included to get the data via FTP and put it into the format needed to initialize the model. The scripts are used in the following way:

data_prep_*model*.com *input1 input2 input3*

model is the form of data to get (AVN, wafs, 104, 212, or 221_tiles).

input1 is the model cycle (00, 06, 12, or 18). It assumes the current day.

input2 is the final forecast hour of model data to get, effectively setting a limit on the length of the workstation Eta forecast.

input3 is the interval in hours between input model forecasts (sets the frequency at which the lateral boundary conditions can be updated; typically 3 or 6).

EXAMPLE: If the user wanted grid 212 data for a 24 hour model run based on the 00Z cycle with boundary condition updates every 3 h, the following command would be entered: **data_prep_212.com** *00 24 03*

2) Special comments about the WAFS form of AVN data, which has global data divided into 8 $90^\circ \times 90^\circ$ octants:

a) The region covered by each octant is shown in the GIF image *wafs_octs.gif* in *worketa_nh/dprep/bin.*

b) The user of this dataset will have to edit the end of *etagrib_wafs.sh* to list the octant numbers to process. This list needs to end with the number –9.

c) WAFS files are small for two main reasons: they contain 11 isobaric levels (the full AVN grids have 26), and the data is "thinned" near the poles to counter the convergence of meridians. The resolution is $1.25^{\circ} \times 1.25^{\circ}$ near the equator.

d) One inconvenience with the WAFS form is that the data is only available at 6 h intervals, decreasing the potential frequency of lateral boundary updates. This negative aspect is offset by file sizes that are about 15-20% as large as the full AVN files–a key factor for low-bandwidth users.

3) Special comments about the grid 221 data, which is obtained in "tiled" form:

a) The region covered by each of the 36 tiles is displayed on the webpage *http://www.emc.ncep.noaa.gov/research/tiles.221.html*. Some tiles along the boundary of this region are not completely filled with data, as this output grid extends beyond the computational domain of the operational Eta along part of the boundary.

b) The user of this dataset will have to edit the files *get_data_221_tiles.com* and *etagrib_221_tile.sh* to obtain the tiles needed for a particular domain. The code which merges these tiles is fairly general, so any normal square/rectangular region should work.

c) The tiled grid 221 dataset provides 32 km horizontal resolution data, and due to the tiling a limited volume of data will cover a typical workstation domain. The region covered includes essentially the entire computational domain of the operational Eta, opening up a vast region to workstation Eta runs.

d) At this time only the most recent model cycle is stored on the server.

4) Surface fields

a) Obtaining the time-dependent fields

sstgrb, sstgrb.index, imssnow.grb, snowdepth.grb, sfcanl

Go to *worketa_nh/eta/bin* and enter **get_sfcfields.com**.

This script uses FTP to bring most recent versions of the snow and SST files into the local directory. The model can run using old snow and SST files, but for better results these should be updated several times a week (particularly if the snowcover is varying over the computational domain).

sfcanl is the surface analysis from an AVN model run; it **must** be used for obtaining soil fields when using WAFS data as the initial dataset, as the WAFS grids themselves do not have soil information. The script that processes WAFS data automatically downloads the appropriate *sfcanl* file.

The logical switch GRIBSOIL in the file *ETAIN* (edited below) specifies whether GRIB-based or *sfcanl*-based soil fields will be used.

b) Descriptions of the included time-independent fields located in *worketa_nh/eta/bin*:

islope_1d_ieee, isltyp_1d_ieee, ivgtyp_1d_ieee:

These specify the slope, soil, and vegetation types.

alb1_ieee, alb2_ieee, alb3_ieee, alb4_ieee:

These are seasonal albedo files which are interpolated in time to get a starting value of albedo. This initial albedo value is refined by code that considers surface characteristics, such as snow/sea-ice cover.

veg.eta.grb:

This file contains the vegetation fraction, stored as a monthly value and interpolated in time to obtain the fraction on the day of the model run.

imsmask.ascii:

This file contains the land/sea mask used when processing the IMS snow data (this is a high-resolution NESDIS product which is used to define snow and ice cover, but not depth).

rfusaflw_ieee:

This file contains the land/sea mask used when processing the U. S. Air Force snow data (a lower-resolution product used to define the model snowdepth).

Setting up and running the model

1) Go to *worketa_nh/eta/bin* and edit the namelist file *ETAIN* to set some details about the model run. NOTE: only edit the *ETAIN* file located in *worketa_nh/eta/bin*; copies in other directories are copied from this location.

TLM0D,TPH0D: These are the center longitude (°E positive, °W negative) and latitude of the model domain.

IM,JM,LM: The dimensions of the computational domain. The IM, JM, and LM values here were set from the values in *set_parmeta* in *worketa_nh/eta/install* and should not be changed.

PTINP: The model top pressure in Pa; the value was previously set in *set_parmeta* and should not be changed here.

DLMD,DPHD: The grid spacing in degrees in the east-west and north-south directions within the rotated lat-lon coordinate of the Eta model. An approximate relationship between horizontal grid spacing and DLMD & DPHD values is given below in the description of DT.

DT: The fundamental time step of the model (in seconds), which depends upon horizontal resolution. Some standard values used at NCEP are given below; they should provide computational stability for all flows. Note that the given values of DT divide evenly into 3600 so the proper number of physics/adjustment steps will be performed each hour.

$DLMD, DPHD$	*appx. hor. res.* (km)	$DT(s)$
.067, .066	10	20
.099, .097	15	30
.154, .141	22	60
.222, .205	32	90
.333, .308	48	120
.577, .538	80	200

An overly long time step leads to model instability, which typically is apparent as a dramatic, non-physical oscillation in surface pressure.

IDTAD: The frequency of advection time steps (should remain 2, which means that advection is called at every other fundamental time step).

IMONTH,IDATE,IYEAR,ISTRTIM: Together these specify the initial time of the run. ISTRTIM is the hour of the starting time (e.g., a run based on 12Z data has ISTRTIM=12)

NSOIL: The number of soil levels (should remain 4).

NINIT: The number of input data files.

INIT_IN(#): The names of the input data files. The naming convention is *YYMMDDCCHHH.ETAmod*, where *YYMMDD* is the year, month and day, *CC* is the cycle time, *HHH* is the hour of the forecast, and *mod* is the extension given to the particular input dataset (104, 212, _avn, _tile, _wafs). It is possible to set up an automated system to automatically specify these file names if running in real-time (see Appendix: Automating the process). The number of files described should equal NINIT specified above.

INIT_OUT: The directory where the initial and boundary condition files will be written (should not need to be changed).

TBOCO: Interval in hours between input model data (typically 3 or 6).

NHOUR: Should be set to $(NINIT-1) \times TBOCO$. This value is the length of time covered by the Eta model boundary condition file, so it represents an upper limit on forecast length.

TOPO_IN: Location where the raw $(10^\circ \times 10^\circ)$ topo files are stored. This was set during a configure step above and should not need to be changed.

TOPO_OUT: File that stores the topography data for the current model domain.

GRIBSOIL: Described above in step #4 of "Getting the input data". If .TRUE. then soil temperature and moisture are initialized from the input GRIB dataset. If .FALSE. then soil fields are initialized from the *sfcanl* file. This *must* be set to .FALSE. if initializing from the WAFS form of the AVN!

seares: The resolution of the land/sea mask data to use (2, 4, or 8). Refer to Step #7 in "Initial setup".

SIGMA: This logical switch defines the vertical coordinate to be used. When SIGMA is set to .false. the Eta coordinate is used; when SIGMA is .true. the Sigma coordinate is used.

2) enter the command: **new_prep.sh**

This script runs programs that generate a topography file and the initial and lateral boundary conditions. The topography needs to be generated only once per domain, so the line of the script running *etatopo.exe* can be commented out if multiple runs are made on the same domain (the output topography is stored in the TOPO_OUT file defined above and is overwritten each time that *etatopo.exe* is run). This script generates two files of standard output, *topo.out* and *initbc.out*. It also generates the data files *etatopo.dat* and *ZEFF* that are used by the interpolation code and the model, respectively.

3) enter the command: **cd ../runs**

4) edit the *fcstdata.meso* file, which controls details of the model run.

```
&FCSTDATA
TSTART=00.0,TEND=24.00,TCP=99.0,RESTRT=.FALSE.,SINGLRST=.TRUE.,SUBPOST=.FALSE.,
NMAP=5,TSHDE=00.0,06.0,12.0,18.0,24.0,99.0,99.0,99.0,99.0,99.0,
              99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,
              99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,
              99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,
              99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,99.0,
SPL=10000.,15000.,20000.,25000.,30000.,35000.,40000.,45000.,50000.,
     55000.,60000.,65000.,70000.,75000.,80000.,85000.,90000.,92500.,
     95000.,100000.
NPHS=8,NCNVC=8,NRADSH=1,NRADLH=2,NTDDMP=1,
TPREC=6.0,THEAT=6.0,TCLOD=6.0,
TRDSW=6.0,TRDLW=6.0,TSRFC=6.0
NEST=.FALSE.,HYDRO=.TRUE.,SPLINE=.TRUE.
&END
```

TSTART: Sets the beginning forecast hour (=0.0 unless restarting a forecast as discussed in Appendix: Restarting a forecast).

TEND: Sets the ending forecast hour.

RESTRT: Logical switch that tells the model whether it is starting from an initial condition file (if =.FALSE.) or from a model restrt file (if =.TRUE.). Typically set to .FALSE.; an example of when it might be true is given in the appendix (Restarting a forecast).

SINGLRST: Should always be .TRUE. for what is described in this manual.

NMAP: Specifies the number of times at which the model will generate output.

TSHDE: List of the forecast hours for which the model will generate output. The model considers the first NMAP values in the list while the list of 99.0 values are space fillers.

SPL: These are the pressure levels (in Pa) at which isobaric output will be generated by the post-processor. The number of levels should equal the LSM value defined in *set_parmeta* in step #9 of the "Initial setup" section.

NPHS: The physics package is called every NPHS time steps. The surface physics are usually run about every 10 minutes of forecast time (NPHS set to an **even** number such that NPHS $\times$ DT $\approx$ 600). For example, a 10 km resolution run using a 20 s time step would set NPHS = 30.

NCNVC: The frequency of calls to the convective parameterization. Should be the same value as NPHS.

NRADSH,NRADLH: The interval in hours between calls to the radiation schemes (shortwave and longwave). These values are typically 1 and 2, respectively.

NTDDMP: The frequency in time steps at which the divergence damping is applied (typically 1).

The next six values specify the number of hours over which the following types of parameters accumulate prior to being reset to zero. They should be set considering the frequency at which output is generated (i.e., resetting precipitation totals every 2 h will be problematic if output is generated every 3 h).

TPREC: precipitation
THEAT: average latent heating associated with precipitation
TCLOD: average cloud fractions
TRDSW: short wave radiation
TRDLW: long wave radiation
TSRFC: surface fluxes (e.g., average sensible heat flux)

NEST: should be set to .FALSE. for workstation Eta purposes.

HYDRO: .TRUE. uses the hydrostatic dynamics, while .FALSE. invokes Janjić's nonhydrostatic extensions to the Eta model. Note: Running in nonhydrostatic mode adds significantly to the run time (in some cases around a 70% increase), and will have its largest impact on forecasts at resolutions finer than 10 km.

SPLINE: Only relevant when running with the sigma coordinate. If .TRUE. a spline fit is used when producing output. This primarily impacts the below-ground parts of the domain, producing much smoother underground fields (such as reduced sea-level pressure). Recommendation is to set the switch to .TRUE.

5) Running the model:

enter the command: **run.com_mpi** possibly with the following option:

> **kf** – will run the model using Kain-Fritsch (KF) convection rather than Betts-Miller-Janjić (BMJ) convection. Due to the very different approaches taken by these convection schemes, significant differences in precipitation and low-level temperatures are possible when switching from one to the other.
>
> If no option is given after the command, BMJ convection will be used.

As the model runs, it will generate restrt files (with names such as *restrt00.quilt.t00s* if running the model on a 1-CPU computer, or restrt00.t00s if the model is run on a greater number of CPUs than INPES × JNPES) at each specified output time. These restrt files contain the model forecast, but must be post-processed to be useful.

The post-processor

1) Run the post-processor with the following command:

outjob_special *stime etime tint type*

where *stime* is the starting time (typically 00), *etime* is the final forecast hour, *tint* is the interval at which restrt files were created during the run.

This script will perform a number of functions: "quilt" the restrt files (primarily calculate a reduced sea-level pressure, which is not done in the integration of the model) if needed (the script will look at the restrt file names and act accordingly), run the post-processor to create GRIB output on the model's native grid, and interpolate the GRIB output onto a standard projection (by setting *type* to either latlon or lmbc).

The remaining discussion in this section is fairly detailed, and not critical for the first-time user. It describes how to control which data gets output by the post (i.e., fields, levels) and how to specify a specific output grid.

Controlling the post

The *cntrl.parm* file (unit 14) controls the fields and levels that are written to the output GRIB file. Each parameter is controlled by two lines in this file.

```
(TEMP ON PRESS SFCS ) Q=( 16), S=( 8), SCAL=(-3.0), SMTH=(00 00 00)
L=(00000 00000 00000 00000 00000 11111 11111 11110 00000 00000 00000 00000)
```

The first line describes the field, sets the Q and S values (no longer used), and defines the precision and smoothing.

SCAL defines the precision that is written out to the GRIB data. Positive values denote decimal scaling (maintain that number of significant digits), while negative values denote binary scaling (precise to 2^{SCAL}; $SCAL = -3.0$ gives output precise to the nearest $\frac{1}{8}$). Higher resolution runs should output the data with more precision if performing complex diagnostics on the data that use higher order derivatives (e.g., Q vectors, vorticity advection).

The final item on the first line is the smoothing block SMTH. The first value sets the number of smoothing passes applied to the staggered (where mass and wind points are separate) grid and the second value sets the number of smoothing passes applied to the filled grid (the post "fills" the native grid so there are mass and wind data at all grid points). The third value activates a 25-point Bleck filter on the output grid, which for the workstation post is the filled grid. The Bleck filter removes small scale features and largely preserves extrema in the data even after multiple applications. As seen in the *cntrl.parm* file, smoothing typically is not performed on most fields.

The second line controls the output levels (and whether a field will be created at all) by providing on/off switches for the output.

For multi-level data the switches (1 = output, 0 = do not output) control which levels to create output for, with the lowest pressure level (furthest from ground) controlled by the leftmost element of the string. This ordering can be reversed by using 2's in place of 1's; 2's control output from the ground up. For isobaric data the levels being switched on and off correspond to the SPL values in the *fcstdata.meso* file, and the first LSM values in the string are used. For single level data, the leftmost number is used as the switch (1 = output, 0 = do not output).

Examples (assuming LSM=39, with SPL values defined every 25 hPa from 50 hPa to 1000 hPa):

```
(MESINGER MEAN SLP ) Q=( 8), S=( 138), SCAL=(-0.1), SMTH=(00 00 00)
L=(10000 00000 00000 00000 00000 00000 00000 00000 00000 00000 00000 00000)

(TEMP ON PRESS SFCS ) Q=( 16), S=( 8), SCAL=(-3.0), SMTH=(00 00 00)
L=(00000 00000 00000 00000 00000 11111 11111 11110 00000 00000 00000 00000)

(TEMP ON PRESS SFCS ) Q=( 16), S=( 8), SCAL=(-3.0), SMTH=(00 00 00)
L=(22222 22222 22220 00000 00000 00000 00000 00000 00000 00000 00000 00000)
```

Translations:

The first output field is a mean sea level pressure reduction designed for use with the Eta model. $SCAL = -0.1$ specifies that the value is stored to a precision of $2^{-0.1}$, or pretty close to the nearest Pa.

The second example is isobaric temperature. The initial 25 0's suppress the output of this field between 50 hPa and 650 hPa, while the subsequent 14 1's generate output between 675 hPa and 1000 hPa. $SCAL = -3.0$ writes out the data to the nearest 0.125.

The third output field is identical to the second, but by using 2's instead of 1's it is clearer that output is being generated for the 14 output pressure levels closest to the ground.

Adding a field

More fields are included in the file *cntrl.example*; the easiest way to add a field is to copy one of these into the *cntrl.parm* file. The file *AVBLFLDS* in *worketa_nh/post* lists all fields available for posting, but in a different format.

Suppose highest freezing level data is desired. Searching through the *AVBLFLDS* file finds the proper two lines.

```
      DATA IFILV(165),AVBL(165),IQ(165),IS(165)
     &                     /1,'HIGHEST FREEZE LVL ',007,204/
```

The crucial piece of information from above is actually the character string in quotes on the second line. The numbers listed after the character string give the GRIB specification of the field. The first value is the GRIB parameter number (007 = geopotential height) and the second is vertical coordinate (204 = highest tropospheric freezing level; a special NCEP definition). The number preceding the character string is a switch that indicates whether the variables are stored on mass (1) or wind (0) grid points.

To add this field into the *cntrl.parm* file, this character string (including any spaces) is placed between parenthesis on the first line of the cntrl file. Proper formatting is crucial in the *cntrl.parm* file; make sure any lines added match the format of the other lines in the file. Q and S can be set to any value as they will be replaced by the GRIB definitions described above as the program runs. Scaling and smoothing are set to their desired values, and since this is single level data, only the leftmost element of the second line is set to 1. The lines added to the *cntrl.parm* file would look something like this:

```
(HIGHEST FREEZE LVL ) Q=( 7), S=( 204), SCAL=(-2.0), SMTH=(00 00 00)
L=(10000 00000 00000 00000 00000 00000 00000 00000 00000 00000 00000 00000)
```

Interpolation to other grid projections

The workstation post creates GRIB output on the native e-grid (files with names like *EGRD3Dhh.t00s*), but this output on the rotated latitude-longitude projection of the Eta model is not readily degribbed by most software packages. The *outjob_special* script uses the included GRIB utility *copygb* to interpolate the native grid data onto regular lat-lon (*latlon_hh*) and Lambert Conic Conformal grids (*lmbc_hh*), where *hh* represents the forecast hour. These interpolated GRIB files on standard map projections should be much easier to degrib.

One item to point out about copygb is that the binary precision specified in the *cntrl.parm* file will not automatically carry over to the copygb output files. The copygb lines within the *outjob_special* script specify a scaling though the "-s" flag, but this scaling is applied equally to all fields. This indiscriminate scaling might make the lmbc and latlon files unnecessarily large. A way to maintain control over the precision of various fields is to use more decimal scaling ($SCAL > 0$) in the *cntrl.parm* file. Decimal scaling from the original EGRD3D file will automatically carry over to the interpolated GRIB fields without using a "-s" option. A possible solution is given in the *cntrl.dec* file which avoids using binary scaling. To use decimal scaling exclusively the following changes would need to be made to the *outjob_special* file: 1) change the symbolic link from *cntrl.parm* to *cntrl.dec*; 2) remove the -s"2,-3" from the copygb.x lines. This purely decimal scaling option was not thoroughly tested, so it may require some tuning by the user.

The post generates the files *outjob_input_lat* and *outjob_input_lmbc* that define the destination grids for copygb. The values in these files are appropriate for the domain and resolution of the model, but they could be modified by the user if desired. (NOTE: The post will overwrite these files each time it runs, so save to another file name if making modifications)

The file *outjob_input_lat* contains a string of numbers beginning with 255 that describes the lat-lon output domain. These numbers have the format:

255 0 IDIM JDIM LAT1 LON1 128 LAT2 LON2 DLAT DLON 64

where

IDIM, JDIM: The dimensions of the output regular lat-lon grid.

LAT1, LON1: The latitude and longitude of the southwest corner of the output grid expressed in millidegrees. For example, a corner point at 15°N, 112.5°W would be expressed as 15000 -112500.

LAT2, LON2: As with LAT1, LON1, but for the northeast corner of the output grid.

DLAT, DLON: These are the spacings between output grid points in millidegrees of latitude and longitude.

The file *outjob_input_lmbc* has a slightly different string of numbers:

255 3 IDIM JDIM LAT1 LON1 8 CLON DX DY 0 64 TLAT1 TLAT2

where

IDIM, JDIM: These are the dimensions of the Lambert Conformal grid.

LAT1, LON1: The southwest corner specified as in the lat-lon output grid.

CLON: The center longitude of the lambert conformal projection. It will default to be the same as TLM0D (the center longitude of the model).

DX,DY: These are the grid spacings at the true latitude in the X and Y directions. Values are expressed in meters, and default to the approximate spacing of the filled native e-grid (with both mass and wind data at all points).

TLAT1, TLAT2: These are the "true latitudes" of the lambert conformal projection. By default they both are equal to the center latitude of the model.

These values should provide sufficient information to "degrib" the data regardless of the software package being used.

Aside for GEMPAK users:

A simple degribbing script for GEMPAK is included in the package. Running *outjob_special* with five arguments instead of four will automatically degrib either the lmbc or latlon data into a GEMPAK datafile in the directory *worketa_nh/gem_out*. The fourth argument is either lmbc or latlon, and the fifth argument is the name of the GEMPAK datafile.

EXAMPLE: To post-process 12 h of model data at a 3 h interval, interpolate to a Lambert Conformal grid, and finally degrib into the GEMPAK file *lmbc_test.grd*, the following command would be used:

outjob_special 00 12 03 lmbc lmbc_test.grd

Appendix

This is a small collection of items that are not relevant to most users, but may be useful for those "pushing the envelope" with the workstation Eta. As a rule they represent items that have not been tested as rigorously as items described in the main part of the manual.

Automating the process

When the workstation Eta is run in real-time over the same domain, the editing process on the *ETAIN* file becomes tedious. After all, the only items being changed are date specifications. A crude sample of how the process can be automated is included with the files *ETAIN_raw* and *new_prep.sh_auto* in the *worketa_nh/eta/bin* directory. The script *new_prep.sh_auto* uses the *sed* command to substitute the current date for variables in the *ETAIN_raw* file, and output an *ETAIN* file. To facilitate running *new_prep.sh_auto* as a crontab job it was required to change directories to an explicit path (equivalent to *worketa_nh/eta/bin*) so the relative paths in the rest of the script would make sense. Similar directory specifications are needed for the "data_prep" job in *worketa_nh/dprep/bin* and with *run.com_mpi* in *worketa_nh/eta/runs*.

The digital filter

The Eta model has a filter that can be applied at the beginning of the forecast to give smoother fields over the first few hours of the forecast. To use it, the call to DIGFILT in *worketa_nh/eta/src/etafcst_all/EBU.F* needs to be uncommented. Then in *worketa_nh/eta/src/etafcst_all/DIGFILT.F* the parameter NTIM needs to be set. NTIM is the number of time steps that the model will integrate forward and backward from the 0 h forecast time while filtering. Larger values of NTIM cut off lower-frequency phenomena and cause more smoothing. Overly large values of NTIM cause the model to fail. A good starting point for NTIM is a value that will give NTIM $\times$ DT $\approx$ 2400 s. Filtering is most useful when going from a coarse initial analysis to a very high resolution Eta model run. Filtered and unfiltered forecasts converge fairly quickly; filtered runs may give more aesthetically pleasing results during the initial 3-6 h of the forecast.

Restarting a forecast

This option would primarily be useful if for some reason the model could not be run over a single block of time. Maybe there is a conflict with other users on a shared computer resource. In any case, one recent addition to the model's capabilities (courtesy of T. Black) is a "true" restart capability. For example, an intended 24 h forecast has only integrated 12 h when the model run must be halted. Rather than starting again from scratch, the run can be restarted at the 12 h forecast point and integrated to completion. The code has the intelligence to skip the necessary records in the lateral boundary condition file. (NOTE: This new capability will create problems if the user tries to restart the model as has been done in the past with multiple lateral boundary condition files).

To use this option, the TSTART value in the *fcstdata.meso* file should be changed to the hour where the model will be restarted, and RESTRT should be set to .TRUE. The values of TSHDE should reflect the output desired from the restarted run. The *run.com_mpi* file needs to be modified in two ways: the removal of restrt* files near the top should be removed or commented out, and the link to *init.file* should be changed to a link to the appropriate restrt file.

Brief list of Eta References

Dynamics/Numerics:

Janjić, Z. I., 1974: A stable centered difference scheme free of two-grid-interval noise. *Mon. Wea. Rev.*, **102**, 319–323.

Janjić, Z. I., 1979: Forward-backward scheme modified to prevent two-grid-interval noise and its application in sigma coordinate models. *Contrib. Atmos. Phys.*, **52**, 69–84.

Janjić, Z. I., 1984: Non-linear advection schemes and energy cascade on semi-staggerd grids. *Mon. Wea. Rev.*, **112**, 1234–1245.

Mesinger, F., 1973: A method for construction of second-order accuracy difference schemes permitting no false two-grid interval wave in the height field. *Tellus*, **25**, 444–458.

Mesinger, F., 1977: Forward-backward scheme, and its use in a limited area model. *Contrib. Atmos. Phys.*, **50**, 200–210.

Van Leer, B., 1977: Towards the ultimate conservative difference scheme: A new approach to numerical convection. *J. Comput. Physics*, **23**, 276–299.

Physics:

Betts, A. K., 1986: A new convective adjustment scheme. Part I: Observational and theoretical basis. *Quart. J. Roy. Meteor. Soc.*, **112**, 677–691.

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.

Betts, A. K., F. Chen, K. E. Mitchell, and Z. I. Janjić, 1997: Assessment of land surface and boundary layer models in two operational versions of the NCEP Eta model using FIFE data. *Mon. Wea. Rev.*, **125**, 2896–2916.

Chen, F., K. Mitchell, J. Schaake, Y. Xue, H.-L. Pan, V. Koren, Q. Y. Duan, M. Ek and A. Betts, 1996: Modeling of land surface evaporation by four schemes and comparison with FIFE observations. *J. Geophy. Research*, **101**, 7251–7268.

Fels, S. B., and M. D. Schwarztkopf, 1975: The simplified exchange approximation: A new method for radiative transfer calculations. *J. Atmos. Sci.*, **32**, 1475–1488.

Janjić, Z. I., 1990: The step-mountain coordinate: physical package. *Mon. Wea. Rev.*, **118**, 1429–1443.

Janjić, 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.

Lacis, A. A., and J. E. Hansen, 1974: A parameterization of the absorption of solar radiation in the earth's atmosphere. *J. Atmos. Sci.*, **31**, 118–133.

Mellor, G. L., and T. Yamada, 1974: A hierarchy of turbulence closure models for planetary boundary layers. *J. Atmos. Sci.*, **31**, 1791–1806.

Mellor, G. L., and T. Yamada, 1982: Development of a turbulence closure model for geophysical fluid problems. *Rev. Geophys. Space Phys.*, **20**, 851–875.

Zhao, Q., and F. H. Carr, 1997: A prognostic cloud scheme for operational NWP models. *Mon. Wea. Rev.*, **125**, 1931–1953.

Zhao, Q., T. L. Black, M. E. Baldwin, 1997: Implementation of the cloud prediction scheme in the Eta model at NCEP. *Wea. Forecasting*, **12**, 697–712.

General:

Black, T., 1994: The new NMC mesoscale Eta model: description and forecast examples. *Wea. Forecasting*, **9**, 265–278.

Mesinger, F., 1984: A blocking technique for representation of mountains in atmospheric models. *Riv. Meteor. Aeronaut.*, **44**, 195–202.

Mesinger, F., and T. L. Black, 1992: On the impact of forecast accuracy of the step mountain (eta) vs. sigma coordinate. *Meteor. Atmos. Phys.*, **50**, 47–60.

Mesinger, F., Z. I. Janjić, S. Nickǒvić, D. Gavrilov and D. G. Deaven, 1988: The step-mountain coordinate: model description and performance for cases of alpine lee cyclogenesis and for a case of an Appalachian redevelopment. *Mon. Wea. Rev.*, **116**, 1493–1518.

Rogers, E., D. G. Deaven, and G. J. DiMego, 1995: The regional analysis system for the operational Eta model: Original 80 km configuration, recent changes, and future plans. *Wea. Forecasting*, **10**, 810–825.

Rogers, E., T. L. Black, D. G. Deaven, G. J. DiMego, and others, 1996: Changes to the operational "early" Eta analysis/forecast system at the National Centers for Environmental Prediction. *Wea. Forecasting*, **11**, 391–413.

Internal NCEP References

Black, T. L., 1988: The step-mountain, Eta coordinate Regional Model: A documentation. NWS/NMC Washington, 47pp [Available from NCEP, 5200 Auth Road, Camp Springs, MD 20746].

Black, T. L., D. Deaven and G. DiMego, 1993: The step-mountain Eta coordinate model: 80-km 'early' version and objective verifications. Technical Procedures Bulletin, No. 412, NOAA/NWS, 31 pp. [Available from National Weather Service, Office of Meteorology, 1325 East-West Highway, Silver Spring, MD 20910]

Dey, C. H., 1996: The WMO format for storage of weather product information and the exchange of weather product messages in gridded binary (GRIB) format. Office Note 388, NOAA/NWS/NCEP. [Available from NCEP, Room 101, 5200 Auth Road, Camp Springs, MD 20746.]

Rogers, E., T. Black, D. Deaven, G. DiMego, Q. Zhao, Y. Lin, N. W. Junker, and M. Baldwin, 1995: Changes to the NMC operational Eta model analysis/forecast system. Technical Procedures Bulletin No. 423, NOAA/NWS, 60 pp. [National Weather Service, Office of Meteorology, 1325 East-West Highway, Silver Spring, MD 20910]

Treadon, R. E., 1993: The NMC Eta Model post processor: A documentation. NMC Office Note 394, NOAA/NWS, 44 pp. [Available from NCEP, Room 101, 5200 Auth Road, Camp Springs, MD 20746.]