

Exercise in Setting Up a Case Study

Goals

- Outline a process for performing a case study
- Provide you with repeated practice - with less hand-holding - setting up and running a WRF simulation
- Provide you with information on input data acquisition from some of the available model datasets.

Scenario (From Eric Stevens' email of 15 May 2009):

Once you get settled in at ARSC, we would really appreciate your help in running the WRF on a case study. From five consecutive hours, from 11z through 15z on February 27 of this year, Savoonga on Saint Lawrence Island had wind gusts of at least 71kt, with a peak wind of 84kt. I suspect that a 2200ft terrain feature, Atuk Mountain, just south of Savoonga played an important role in locally modifying the broad-scale pattern to produce a local acceleration of the winds at Savoonga, but this is just a conjecture without objective support so far. Since the terrain feature in question is comparatively subtle, perhaps we need to use a very-fine resolution in the WRF to resolve the phenomena. Can we go to 500 or 250 meters with the WRF? I think Carl Dierking in Juneau did something like this in his area with the WRF. If we initialized the model with the 40km GFS forcing the lateral boundary conditions, how much time do you think would be required for the model to scale down to the details? Could we run the WRF based on the 06Z GFS? The good news is that we only need to stretch the innermost nest over a small area (Saint Lawrence Island), so that should save on crunching time.

Considerations

- Not sure if the 40km GFS is available in an archive.
- We could initialize with a 1.0 degree (approximately 100km) [NCEP FNL \(Final\) Operational Global Analysis](#) (from GFS). Since this is global, we wouldn't have to worry about our domain falling outside of the input domain. It's fairly easy to ingest GFS into WPS.
- [North American Regional Reanalysis \(NARR\)](#) is available at approximately 32km resolution, but has a more limited domain [\[Graphic of NARR Domain\]](#). Also, NARR requires more work to ingest into WRF.
- Because I happen to have the NAM AWIPAK (45km resolution) forecast for the desired time period still stored in `/projects/ARSCWTHR/InputData/AWIPAK`, I could use this. It's easy to ingest into WPS, but the domain is regional, so we need to make sure that our desired domain will completely fit into the AWIPAK domain.

The Game Plan

- Since it's simpler, we'll start by using FNL for initial and lateral boundary conditions.
- We'll initialize the model at 2009-02-27_00Z, which will provide an 11-hour "spinup" to the wind event.
- For a first attempt, we'll do a single 50x50 grid point nest at 27 km resolution, centered roughly on Saint Lawrence Island. By consulting a map, I find a lower-left corner for a bounding box on Saint Lawrence Island is roughly 63N, 172W, and an upper-left corner is at roughly 64N, 168W. So, I'll choose a center lat/lon of 63.5N / 170.0W.

The Steps

Run the *userWRFSetup.sh* script to set up directories for your case study

Use the FNL for initialization. I have copies of the necessary files in

/archive/u1/uaf/morton/InputData/FNL/2009-02/

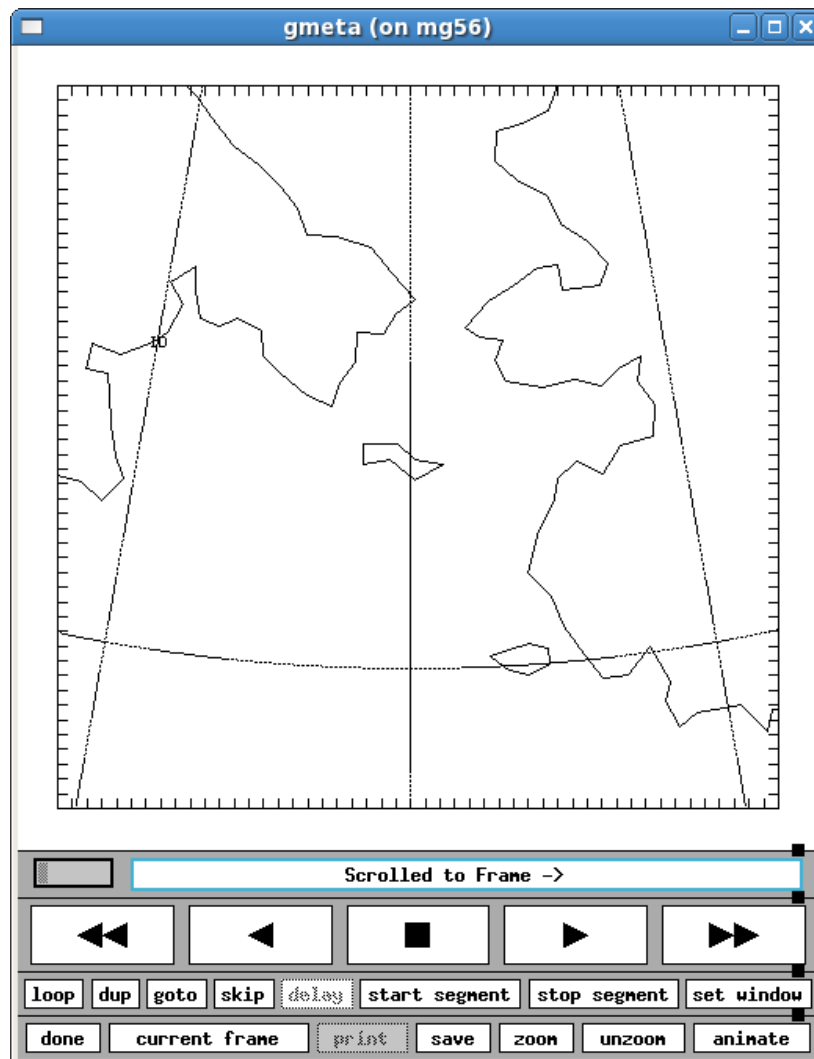
For your future reference, this is how I got the files:

- Go to <http://dss.ucar.edu/datasets/ds083.2/> and sign-in (or register for username/password, if necessary)
- Click on the *Data Access* tab, then the *Internet Download* link
- Click on the *tabular Internet Download file list* link
- Select the *Group ID* (year) of interest, which is 2009 in our case
- Scroll down to the file list for 2009.02
- In this list, select the files for the desired time period:
 - *fnl_090227_00_00_c*
 - *fnl_090227_06_00_c*
 - *fnl_090227_12_00_c*
 - *fnl_090227_18_00_c*
 - *fnl_090228_00_00_c*
- Retrieve them, using one of the available methods

Create the domain

- Decide on region of interest, resolution, center point, etc. For this case, I decided on
 - Center on Saint Lawrence Island (approximately 63.5N, 170.0W)

- Since I'm using GFS for input, I don't need to worry about restricting my domain, so I'm going with a 50x50 grid at 27km resolution. I chose 27km, because this would be easy to nest down to 9km, 3km and 1km resolution.
- Modify *namelist.wps* appropriately
- Run *util/plotgrids.exe*, then view the resulting *gmeta* file to see if the domain looks reasonable, and to insure that you didn't accidentally set this up for an Antarctic case study. The following image was obtained by using *idt* to view the *gmeta* file produced by *plotgrids.exe* on a 50x50 grid with 27km resolution centered on Saint Lawrence Island.



- When satisfied, run *geogrid.exe*
- Confirm presence of *geo_em.d01.nc* and use *ncview* to verify

Ingest the input data - the idea is to convert the input data from a native GRIB format to an intermediate format understood by WRF preprocessing routines.

- Be sure that the *namelist.wps* dates and intervals are set up correctly.
- Make a link to the correct Vtable (for this FNL data, it would be *Vtable.GFS*)
- Use *link_grib.csh* to make links to the input files
- Run *ungrib.exe*
- Confirm presence of the ungribbed, intermediate files
- For your information, there are several utilities available for checking the original (GRIB) and intermediate-formatted input data
 - *util/g1print.exe filename* - Prints a listing of the fields in a GRIB1 file
 - *util/g2print.exe filename* - Prints a listing of the fields in a GRIB2 file
 - *util/plotfmt.exe filename* - Creates a *gmeta* file plot (viewable with *idt*) of the contents of the intermediate format file
 - *util/rd_intermediate.exe filename* - Prints information about the fields in an intermediate format file.

Run *metgrid.exe* to produce files (*met_em**) which will contain the input data mapped to your model grid. These are netCDF files, so may be reviewed with a netCDF viewing tool such as *ncview*.

Change directory over to *../WRFRun* to produce the initial and lateral boundary condition files (*wrfinput_d01* and *wrfbdy_d01*) with *real.exe*, and then to run WRF with *wrf.exe*.

- Set up *namelist.input*
 - you should fill in the necessary fields in the *&time_control* and *&domains* sections. Remember to get the correct value for *num_metgrid_levels* by using *ncdump* on one of the *met_em** files.
 - for the *time_step* value (in seconds), a rough rule of thumb is to set it to less than three times the resolution (in km). If you set it much larger, you may run into model instabilities, and if you set it too small, the simulation will take much longer to run - it's a tradeoff. In this case, with a resolution of 27km, I would be inclined to set it at 90s. However, given that this is a high-wind event, subject to instabilities, I'm even more inclined to set it at something smaller, like 60s.

- real-serial.exe
 - Copy or make a link to the *met_em** files in *../WPS*, then run *real-serial.exe*.
 - If this is a huge domain, you should consider running *real.exe* (the parallel version) by editing and submitting the *realrun.pbs* PBS script.
 - Verify that new *wrfinput_d01* and *wrfbdy_d01* input files were created. These are netCDF files and may be further scrutinized with netCDF viewing tools.
- wrf.exe
 - If this is a small case, and/or you just want to verify that it will run for a few timesteps, you can run directly from the command line using the *wrf-serial.exe* executable.
 - **ARSC DISCOURAGES USERS FROM RUNNING COMPUTATIONALLY INTENSIVE JOBS FROM THE COMMAND LINE. THIS IS WHAT THE COMPUTE NODES ARE FOR.**
 - If your simulation involves more than about 50x50 grid points in a single nest and/or you expect the simulation to take more than about 10 minutes, then you should think of running your simulation on multiple cpu's by editing and submitting *wrfun.pbs*
- The output
 - The output for this particular case will go to the netCDF file, *wrfout_d01_2009-02-27_00:00:00*. If you use *ncview* for a quick look at the output, I recommend selecting the *U10* and *V10* (10 meter wind) variables from the *3d vars* button in *ncview* (the three dimensions are *x*, *y* and *t*). By using the "play" controls, you can timestep through the output. I found that you can see an interesting "surge" of high wind move across the model domain, and it might be interesting to run the model again, creating output at every 30 to 60 minutes, rather than every 180 minutes.

Other things you might try, depending on your time and interest

- Start another similar case study that has a 9km nest inserted. But, why stop there? Why not try a 3km nest inside of the 9km nest, and then a 1km nest inside the 3km nest?!
- There are other input datasets available (described in the Appendix of this document). Try to run the same case study with different input data
- Dream up your own case study using a different region and/or time. You can roughly follow the above procedures just to keep reminding yourself of the general work flow.

Appendix A - Using North American Regional Reanalysis

It can take hours to download all the necessary NARR input files, so I have pre-fetched the data for this case study and made it available in

/archive/u1/uaf/morton/InputData/NARR/2009-02

For your future reference, here are the steps I used to acquire the data:

- Sign in (or register for a userid/password first) to <http://dss.ucar.edu/> (sign-in/register is in the upper left corner).
- Go to the Quick Link for the NARR Reanalyses, or go directly to <http://dss.ucar.edu/pub/narr/>
- Click on the "3-hourly Data Files" link on the left
- Click on the 200902 (February 2009) link in the main page to access the files available for that month/year.
- Note that filenames are of the form **NARRxx_YYYYMM_DDDD.tar**, where **xx** refers to variable class (3D, flx, sfc, etc.) and **DDDD** refer to the start day and end day of the data in the tar file. We need the 3D and sfc files that correspond to our times of interest. In some case studies, you'll find that your days of interest are contained across multiple files, and you may have to end up downloading a lot of data just because your time of interest was on the boundaries of the files. We need the data for just one day - 27 February 2009, so will retrieve the following files. Note that we need to download two 3D files, because the data for 00Z on 28 February is located in the second file:
 - NARR3D_200902_2527.tar (3D fields)
 - NARR3D_200902_2828.tar (3D fields)
 - NARRsfc_200902_2028.tar (surface fields)
- There are a number of ways to retrieve this data, as described on the web page. The details of access are beyond the scope of this particular tutorial. However, the way I like to do it is to check the files I'm interested in, click on the "Create a Wget Script" button, then cut and paste the resulting script statements into a script on the machine where I want to store the data (e.g. the \$ARCHIVE directory of midnight or pingo). Once the wget script is in place, the only thing left to do is to fill in the password in the appropriate place, then run the script.
- You also need to access a file with the "fixed" fields - this contains static fields, like topography, so only needs to be retrieved once. You

can retrieve this by going back to the main NARR selection window at <http://dss.ucar.edu/pub/narr/>, selecting "NARR Fixed Fields" (under the 3-hourly Data Files link), then retrieving the 32km output grid GRIB file (don't grab the *32km native grid GRIB file*). Or, I have a copy, accessible from midnight and pingo, at

/archive/u1/uaf/morton/InputData/NARR/FixedFields/32km_output.AWIP32.fixed

Set up a case study directory

Set up the domain

- Make sure we create our model domain so that it fits in the NARR domain.
- Edit *namelist.wps*, verify correct boundaries by running *util/plotgrids.exe*, then run *geogrid.exe*
- Verify fields in *geo_em.d01.nc*

And now, for the hard part. The three types of NARR files (3D, sfc and fixed) need to be ungribbed separately, and then prior to running *metgrid.exe* we need to modify *namelist.wps* to indicate the multiple types of files that will be used to produce the *met_em** output of *metgrid.exe*. Although this procedure has been outlined in the *WRF Input Data* PPT slides for this tutorial, I am outlining the process below:

- Make the link *Vtable* to *ungrib/Variable_Tables/Vtable.NARR*
- Set up the correct start/stop time and interval (10800 seconds) in *namelist.wps*
- Ungrib the NARR files
 - In *namelist.wps*, in the *&ungrib* section, change

prefix = 'FILE',

to

prefix = 'NARR',

- You can actually set prefix to anything you want - it just indicates the prefix label that will be applied to the ungribbed intermediate format files.
- Run *link_grib.csh*, giving it as an argument the list of NARR files to ungrib

./link_grib.csh /archive/u1/uaf/morton/InputData/NARR/2009-02/

NARR*

- Run `./ungrib.exe` and verify creation of ungribbed intermediate format files in three-hour increments from **NARR:2009-02-27_00** to **NARR:2009-02-28_00**.
- Remember, you can use `util/plotfmt.exe` to plot one of the ungribbed files to see if things look right.
- Ungrib the fixed-field file - pay close attention here. The fixed field file contains constants (like topography) with a single timestamp of 00Z on 08 November 1979. To ungrib this file, you will have to change the start/stop date in `namelist.wps` to 1979-11-08_00:00:00, run `ungrib.exe`, then, before running `metgrid.exe`, change the start/stop times in `namelist.wps` back to their original values. Because this is tedious and error prone, I often like to save a copy of `namelist.wps` before changing the start/stop time for the fixed field, then after I've ungribbed the fixed field file, I just go back to using the copy I saved. So, the steps are
 - In `namelist.wps`, in the `&ungrib` section, set

prefix = 'NARR_FIXED',

- In `namelist.wps`, set the start/stop times to 1979-11-08_00:00:00
- Run `link_grib.csh`, giving it as an argument the fixed field file to ungrib

./link_grib.csh /archive/u1/uaf/morton/InputData/NARR/

FixedFields/*

- Run `./ungrib.exe` and verify creation of an ungribbed intermediate format file in three-hour increments from `NARR_FIXED:1979-11-08_00`.
- Remember to change the start/stop time in `namelist.wps` back to the original values before running `metgrid.exe`!
- Run `metgrid` to produce the `met_em*` files containing the input data mapped to the grid you defined for this case study
 - Again, be sure that the start/stop times in `namelist.wps` are correct.
 - In `namelist.wps`, edit the `&metgrid` section as follows. The idea is that you want to tell `metgrid` to get its input from both the `NARR_FIXED` intermediate file and the `NARR` files:

&metgrid

constants_name = 'NARR_FIXED:1979-11-08_00'

fg_name = 'NARR'

```
io_form_metgrid = 2,  
/
```

- Run *./metgrid.exe* and verify creation of the *met_em** files
- Particularly in a case like this where we have our domain set up near a boundary of the input domain, it's important to be sure that we have input data for all of our model points. Using *ncview* to take a quick look at some of the fields in one of the *met_em** files can be valuable here. Experience suggests that if we have missing input data (due to the fact that our domain might extend beyond the boundaries of the input data), there will be no explicit error messages, and the errors will manifest themselves somewhere downstream in the process, and it won't be readily apparent that the problems are coming from missing input data.
- From this point on, you proceed ahead as you did before. The difficult part here has been to ungrib input files from different sources, then combine their data into the *met_em** input files mapped to your specific domain. This approach has also been used to introduce better sea-ice data into a model, where the sea-ice representation in the atmospheric model (e.g. GFS or NAM) may have been deficient.

real.exe - in the *../WRFRun* directory

- Make links to the *met_em** files in *../WPS*
- Use *ncdump* on one of the *met_em** files to find *num_metgrid_levels* (it's 30, for this case)
- Edit *namelist.input* as before, but, additionally add the line *p_top_requested = 10000* into the *&domains* section. By default, WRF will assume a top of atmosphere at 5000 Pa, but the NARR input data only goes up to 10000 Pa. So, this statement tells WRF to use this as the top of atmosphere.
- Run *real-serial.exe* (or, for a huge domain, submit a parallel job with the *realrun.pbs* PBS script) and verify creation of the initial conditions file, *wrfinput_d01* and the lateral boundary conditions file, *wrfbdy_d01*.

Finally, run *wrf.exe* as you normally would.

Appendix B - Using Real-time input data

The AWIPAK data a 45km grid over the Alaska region, derived from the NAM forecast. I've stored the data needed for this simulation in

/archive/u1/uaf/morton/InputData/AWIPAK/2009022700/

Last Updated 19 June 2009, 23:09Z

The current AWIPAK forecasts are stored in

/projects/ARSCWTHR/InputData/AWIPAK/

You should always be able to find current forecasts here, but they get purged after a while.

Likewise, current GFS 1.0 degree forecast grids are available in

/projects/ARSCWTHR/InputData/GFS1p0

and these are also purged after a while.