In the next few sections, we will look at three different ways to configure HYSPLIT to do dust emissions during a dust storm. In the first, we will look at the constant threshold method. This is the original procedure that was set up for applications in Asia, the Middle East, and the Sahara. The newest procedure, we will discuss in the subsequent sections, was developed for use over North America. For both of these versions, the dust emission algorithm is invoked by setting the name list parameter ICHEM=3. In addition, the CONTROL file requires two more variables after the starting location and height. We will discuss to discuss this shortly.

The original algorithm, which is not currently available, describes the vertical dust flux as a fraction of the horizontal transport and the meteorological variable was the friction velocity, that is another measure of how much momentum is transported to the surface by the wind. The greater the momentum transport, the greater the dust flux from the surface. The important factor was the exceedance of the friction velocity over a threshold friction velocity. And this threshold friction velocity defined, it was either on or off, so that if the friction velocity was less than the threshold, there would be no emissions. And this constant approach, the simplified algorithm, we assume a constant value for the threshold friction velocity, and that was 28 cm/s, and it applied to the desert land use category. So this equation was simplified to this, so that it was only dependent upon the friction velocity and an area, where the friction velocity needed to exceed that constant threshold for dust emissions.

And the approach with HYSPLIT works the same way in the, or the way we use the matrix configuration, where we define 3 starting points. If you recall, we would define a lower left and an upper right corner, and then an increment for the first grid point inside of that box to define the resolution. And what the preprocessor will do is read that file and find all the desert land use points within that domain. And that what's invoked when you do a special run for dust storms. However the problem is in the United States, and that's why this algorithm was not applied originally in the United States, was that the 1° land-use files which are the default that is supplied with HYSPLIT, have no desert land use categories in the Continental United States. So the solution to this is to go into the HYSPLIT4 directory in the bdyfiles directory, and in this directory is contained another subdirectory called bdyfiles0p5 or 0.5. And this contains half-degree resolution land-use, roughness length, and terrain. So to use this file instead of the default one degree, what we would do is cut/copy this file, the ASCDATA.CFG file, and put it in the hysplit4/working directory, and replacing the file currently in that destination. And if you were to look at this file in Notepad, you can see that it changes the pointer, and now looks for the land use and terrain and roughness length in the bdyfiles0p5 directory, and this is a half-degree resolution data file. So that’s all you basically need to do to get the algorithm to work over the continental U.S.

The next step, and I should say that the files that you need to configure the example that we will be doing for dust, can be found in the tutorial/dust directory. So everything we need to configure the run and do the simulation can be found here. And we will go through the steps. So the first thing, let's load the CONTROL file for this simulation, and I should say that we will be doing a simulation for an event, a dust storm, a real dust storm event that occurred in Utah at the end of March 2010. And we do have measurement data, the AirNow network measurement data, so we can compare our model predictions with the measurements.

So to configure this case let's start with the CONTROL file and we will retrieve that from the dust directory. So tutorial, dust, and this first example that we are doing is control_dust0.txt. I'll review the settings in just a moment. Let's also retrieve the name list file and that would be the setup file. Here we go. Like I said, we will review this in a moment. So let’s go back to the CONTROL file and you can see that we’re doing this simulation starting on March 30 of 2010 at 00 hour UTC, and we’re defining a domain about around Utah, around the Great Salt Lake, 3° in latitude and 5° in longitude. This is the lower left and upper right and we're going to set up an interior emission grid of .25, or a quarter degree resolution. And we will do the simulation for 30 hours using the NAM model data for that period. If you are doing this simulation through the web you can find this at the ARL website. As far as the other settings here, we have set the character ID as PM10, and we’re going to set an emission rate of zero, because unless the emission algorithm indicates that we want emissions, we do not want to have emissions, and we're going to cover the entire period of the simulation.   And we're setting up a relatively high resolution concentration grid of .05° with a 30° span in both latitude and longitude, and we’re going to name this output file dust zero to indicate that it is for the initial configuration that we're doing.   And we will have one concentration, air concentration level, which will start collecting, start the model sampling, 6 hours after the start of the simulation and we will collect a 24 hour duration sample. The reason it starts 6 hours after is to match up with the AirNow monitoring network for which we have daily samples for.   And as far as deposition is concerned, we're not going to do deposition right now. And we can save.

And let's take a look at the, at the name list file, and as I noted in the introduction, we enabled the dust algorithm in menu number 10, by selecting this, and sets the ICHEM parameter to three. And we also need to ensure that we are releasing a sufficient number of particles, and that would be menu number four.   And notice what we’re doing here is that we don't need to set the duration, first of all the 24 hours, we can let the particles stay in the computation, but note that we're actually setting the particle release rate per cycle at 100,000 while the maximum number of particles is only 50,000.   So you might think that you should not be able to exceed, in terms of particle release rate, the maximum number of particles. But in fact these are really two different parameters.   The maximum number of particles represents, simply the maximum number of particles that are permitted at any one time during the simulation, because it represents the array space that was allocated to the particle variables. However, the particles released per cycle is simply that, how many particles get released per emission point per emission time.   And the algorithm to compute this, for instance if we have one source and we had an instantaneous release over one time step, like we did in some of the introductory sections, then that hundred thousand particles would be released immediately, and we of course, immediately would be exceeding that limit.   But in this case we're defining a grid of multiple sources with emissions occurring over multiple time periods. Therefore the maximum particle release rate per source is really the hundred thousand divided by the number of sources, by the number of the emission times.   So in fact we could be only releasing a few particles per source whenever that source happens to be emitting and a total may never even get close to 50,000. That's why that was set up that way.

So at this point you're ready to do the simulation, so as you might guess, it is a special run, and it's a dust storm simulation, which means that the program is calling the pre-processor, that is you will get this message, that we have three starting locations.   And this pre-processor does several things.   And let's see what happened here. I see, so we had things left over from the last simulation. So I'm going to go back to this menu, do a reset then a retrieve, and now we should be able to do run model.   So we had some left over things from the smoke simulation, that is the emissions file.

Now while this is running, let me just take a look in the working directory. And first of all you can see the original CONTROL file that we configured, which is under default_conc here with the three starting locations. And, well the two corners of the grid and the resolution, and then the preprocessor created the CONTROL file with 26 starting locations. Now as I mentioned, dust simulations have a special, two additional fields in the CONTROL file, and in this particular case for the original algorithm, the simplified algorithm, the first field is always -1, and then the second field represents the area of the grid cell in square meters, the area of the emission grid cell. So the emission algorithm predicts the emission rate in grams per square meter, this is multiplied then by the area to get the grams emitted, in the desert, added to the particles. Other than that you can see the CONTROL file is the same. And we're getting toward the end, yes. And the last thing of course is the name list file that was created, and you can see the ICHEM is 3, which defines the dust emission algorithm. Other than that everything else is default.

The reason, the first-time, the simulation failed, is because we had defined this emission file for smoke that was left over. When we did a retrieve of the name list, it only overrode those variables that were in the file that we retrieved. So if we had a name list variable within the graphical user interface that was set, that was not in the retrieve, then they would not have been changed. This may or may not be considered a flaw, but what I should have done was do a reset prior to doing the retrieve, at least for the name list. I don't believe it is necessary for the CONTROL file.

Now once the simulation completes, we will look at the measurement data. And the measurement data for that is contained in the tutorial/dust directory. This is the AirNow data for several locations. This is airnow_dust and this is in the DATEM format, and these are the PM10 values in micrograms per cubic meter. And the sample durations are all 24 hours, and they start at 0600. These are the different station IDs that surround the Salt Lake City area. The simulation is almost complete.

And now we can go ahead and do the display and we need to convert to micrograms from grams. Therefore, there are 10^6 micrograms per gram, and we will use these same contour intervals as we used for smoke, and then we can execute display. You can see that the different, so not all of the grid cells within the domain emitted dust, only the grid cells, the emission grid cells at quarter degree resolution here, these are all a quarter of a degree apart, only the ones that the threshold friction velocity exceeded …, well first of all, these are the grid cells with a desert land use category and then to emit, the threshold friction velocity of 28 cm/s centimeters per second had to be exceeded.

We could also plot the measured data on top of that, and let’s see if that works, I did not try this before. Here we go. You can see here, these are where the AirNow measurements occurred, these are the locations downwind, and there are some high values, very much in the hundreds, and we are in the green region, which is, which was only in the tens to 50 or so micrograms.

So let's actually do a statistical analysis, so we can go utility, Convert to DATEM, and we'll set the, we will this here, because it is the same field as in the concentration plotting program. And we need to enter the field here and create the DATEM file, and we can compute statistics, and I'm going to quit that for a second, and try that again. We need to set the contingency level. Took a few tries.   What we should've done was to made sure these were all set and the way to do the was to do a reset when we started the graphical user interface for the first time.

In any event what I wanted to show you, this value here, the average ratio of calculated to measured, is that we're under predicting by a factor of 10, which you saw actually from the graph as well, where the numbers were in the 50s rather than 500.   That would be confirmed by the scatter plot, where most of the predictions, were all on this factor of 10 line of under-prediction.

So, the trick here was that the problem with this menu was that it was not fully populated with the fields; some of them were undefined, that's why all these were originally checked.   So you need to make sure that everything was properly defined and it was not because of the way we continued on from the dust simulation, the smoke simulation, things needed to be reset.

In any event this concludes the simplified dust model configuration and simulation.