Up to this point the only output we have examined from the concentration simulations were air concentration contours. These represent the mass of the particles summed to the concentration grid divided by the volume of the grid cell. However, there is an option to also examine the individual particle positions. This could be useful for diagnostic purposes or to understand the properties of the flow that is controlling the dispersal of the particles. The particle views are three-dimensional, whereas the air concentration contour patterns are on a planar surface.

If you're continuing from the previous section, then the variables in the GUI should already be loaded with the proper values. If not, retrieve the CONTROL file and the name list file that you had created previously, or if you did not create them, they can be retrieved from the tutorial/files directory. So open up the set up run menu and we will do the retrieve, even though we have the correct configuration.

Now for this simulation we want to output particle positions. So we're not that interested in the resolution of the concentration grid. To speed up the calculation, go ahead and make the grid courser. So open up the grids menu and change the spacing back to 5 km, 0.05°, and then save.

The next thing we need to do is set one of the name list parameters to tell the model to output particle positions. In the advanced configuration concentration menu, in menu number nine, you will be outputting a file called

PARDUMP, after three hours and then every three hours thereafter. The default is to have these values at zero. There is no particle file output.   Save, and save, and now run model, using the setup file.

When it completes go to display, but this time particle. And the menu's already pre-populated with the binary input file of particle positions, PARDUMP, the output file will be called partplot.ps, of course.   And the default particle view is cross-section.   You can try the other options on your own later.   Now we know that 10,000 particles are released.   We can make this plot generate a little faster by only plotting every 5th particle.

So go ahead and execute display now.   And what you see here is a projection in the vertical and this projection is along the dotted line.   So the dotted line is drawn automatically through the bulk of the particle positions. And this is after three hours and you can see the particles have mixed up to almost two km above ground level.   You can see here we're plotting 2000 particles, we skipped every …, every fifth one is being plotted.   Going on to six hours later, you can see now we've got particles mixing a few up to almost 3 km.   At nine hours, and finally at 12 hours.   So you can see from this illustration several things.   One, the particles that are farther aloft, are farther downwind.   We already know that wind speeds increase with height above the ground.   And also wind directions tend to veer with height, that is are more clockwise. And so the upper level trajectories also tend to be further south of the, let's say, mean particle pathway.

Now just as a summary, when we talk about particles, we are not talking about physical particles, we're talking about computational particles. These particles that you see represented here, the particle trajectories, are just representations of something else, it could be a gas or it could be a particle, a dust particle. So those characteristics are set in the menu later on. If we were trying to track, for instance, dust particles, we would define a density and a size, and the model would compute a gravitational settling velocity. But for now we're just dealing with these ideal computational particles.

And that concludes the particle display discussion.