In section 8 we will configure HYSPLIT to actually produce realistic concentration output by configuring the model to simulate CAPTEX release number two.

To start with this, let's click on the reset button, and looking to reconfigure the simulation going through each step in sequence. So start with the set up run menu and we're going to start the simulation in 83 09 25 17. We've made these entries before. We've made these entries before. We will select the Dayton, Ohio, release location. We will run for the duration of the sampling, which is 68 hours; naturally in the forward direction. We will use the North American Regional Reanalysis data.

And we're going to now configure the pollutant characteristics. So let's start with the pollutant and we will give this a four character identification which is going to correspond with the actual tracer that was used, PMCH, for perfluro-monomethylcyclohexane. It is an inert non-depositing chemical. You can go to the CAPTEX experiment report and find out that for release number two, the emission rate was 67,000 grams per hour, and the release duration was three hours. In HYSPLIT when time units are defined as zero, that means that it falls back to the default, in this case the time that was defined to start the simulation. So the release actually started at 17 UTC.

Let's configure the concentration. Now we already know that the tracer went from west to east, so instead of centering the concentration grid over the source point,

which is what happens when these numbers are zero, let’s center it somewhat further west, and a little bit to the north, perhaps 42 N and 78 W, would make a good location. And let's make a relatively coarse concentration grid for now, so these simulations will run quickly. Let's make it 25 km, 0.25. We could leave this at 30 or we can make it a little smaller, it makes some of the output files a little smaller, maybe 15 and 25. The plume is spread out more, the grid is spread out more east-west, rather than north-south.

And we will write to the working directory, that's the default, and let’s give it a unique name for the output, let's call it the hysplit output for release number two, and just so it's distinguished from other files, this is a binary file, so we will append the .bin.

Now we will define, only for now, only one vertical level, the depth of the concentration grid cell, and let's leave the default of 100 m. Now, at this point you might be thinking, well if you read the report correctly, the sampling occurred perhaps at 2 m above ground, where the intake for these sampling instruments. So why should I not set this to 2 m, why 100? And we've actually went over the reason for this, and that is as the concentration grid becomes finer, so this is true not just in the horizontal, but also into vertical, you still need to sample a sufficient number of particles to get a realistic simulation, or real realistic view of the distribution of particles. If the concentration grid is very small, you would need many more particles to properly sample the distribution. So it's

always a compromise between resolution, and let's say accuracy, fidelity, and speed. The hundred meters is a good compromise because the vertical resolution of the meteorological data is on the order of hundred meters. So even if we defined more levels, we would not really be resolving any more complex flow features or plume features.

And for the standpoint of this sampling, if you go to the report, you'll find that the sampling actually started, not at the beginning of the simulation, 1700, but actually started at 1800 UTC. So we will explicitly put in the start of the sampling, because we want the model output, the start and stop times of each sampling period, to line up with the sampling data, the measured data, and the sampling continued through the next three days, basically, not quite 72 hours. And we're going to output the sampling interval, the averaging time as three hours. The samples that were collected were either of three-hour duration or of six hour duration. By outputting concentration at three hour intervals we can always put together the six hour average equivalent. Whereas if we output every six hours, as six-hour averages, we could not reverse those averages to come up with the equivalent three hour numbers. So we have to output the finest resolution that was collected.

Now the tracer was inert, it does not deposit, so all these numbers are zero. We will do some experiments later on with some other depositing pollutants. The model is capable of running with multiple pollutants at the same

time, we will discuss that later on. At this point, we only define one tracer, one pollutant. If you try to select something else, it is going to give you a warning message.

So let's close this, and the next thing is, now we know that the default particle number release rate is 2500 over the release cycle. So that, what that means is over the three hour emissions, we would release the particles at an interval, such that after three hours, we had 2500 particles, which would contain approximately 200,000 grams of the perfluorocarbon, of the tracer. This is too small a number for doing a three day duration simulation, so we should up this number, and for now let's use 50,000. We will experiment with this later on as to why this is optimal. And also let's increase the maximum to 100,000. So there's no penalty here for doing this right now but it does allocate a little more array space.

And for diagnostic purposes let's also save the particle output. And let's line up the particle output with the sampling data. So we know that the sampling collection starts at 18, 1800 hrs, the model simulation starts at 17, so we would want the first output to occur one hour after the start of the simulation, and then let's repeat that output every three hours.

Now we just configured HYSPLIT for the perfluorocarbon tracer release. Now nowhere did I enter the units of grams. The model, when it computes the dispersal of the particles, it does not care about units in this context. It only cares about units if we're going to do something with

that mass. If we want to convert to volume units, liters per liter, rather than mass units, grams per liter, so we do those conversions, then it's important for the model to know what units it's working with. But in this case if we enter grams for the input, we will get grams per cubic meter for output. If we had entered kilograms rather than grams as the input, instead of 67,000, we entered 67 as the release rate, 67 units per hour, 67 kg per hour, the output would then be in kilograms per cubic meter. So that's implicit, it's never explicitly set in HYSPLIT. And you as the user have to know that that's what you did.

So let's see if this works. We can just go ahead and do concentration, run model, with that set up file. In the next section we will go over all the output options that are available to us once the simulation has completed.

Once all the particles have been released, after the first three hours, the computational speed is about the same, toward the end of the simulation, the computational speed will appear to be faster because particles are exiting the domain.

And so let's confirm that this simulation was correct. Opening the MESSAGE file, and you can see in the first three hours, that will be this period, this is the end of three hours the emissions have terminated here. We have a little bit over 50,000 particles, we don't exactly get 50,000 because it has to partition the particle release rate according to the time step and we have approximate 201 kg, over 200,000 g of tracer that was released. And you

can see after six hours, the tracer mixed up to about, most of the mass is below 1500 m, which we already know is approximately in depth of the mixed layer, with a few of the particles exceeding that depth. And after 12 hours the distribution is going higher. And I'm just going to scroll down here, until approximately 30 hours into the simulation, and now you can see the particle number starting to go down, and the reason the particle number goes down is that particles are exiting the computational domain, the meteorological data domain. Remember this is not a global data set. This is an extract and it just covers the northeast United States. And by the end of the simulation, we've reduced the particle number to half. So that is the reason why it is running much faster. If I were to have done this calculation using global data, global meteorological data, none of the particles would've exited, and we would've had a 50,000 particle number right to the end. But in a sense, we would be following particles that we don't care about, because once they exit the sampling domain, unless we have some reason to think that they might curve back, then it's computationally efficient for us to just focus our calculations on the area of interest.

And that concludes the configuration of the model for CAPTEX release number two. In the next section we will see what these results look like using different kinds of viewing options.