PUFF SPLIT-MERGE ISSUES

The namelist parameters KSPL, FRHS, FRVS, FRTS, KRND, FRMR, and FRME control the split-merge routines. Normally these should all be left at their default values. Split routines are called at KSPL intervals and merging is always called hourly. Merging is most sensitive to the horizontal parameter FRHS. When going from the default value of 1 to 4 almost all the puffs are merged once FRHS=4. Because merging is called after splitting, most puffs that are merged are already in the same vertical position; hence there is little sensitivity to the vertical parameter, FRVS. The time parameter (FRTS) only matters when there are continuous emissions. KRND controls the interval at which the enhanced merging routines are invoked. Enhanced merging is similar to standard merging except the parameters are 50% larger and selectively applied to those puffs at the lower end of the mass range as defined by FRME. Its default value of 0.10 means that only puffs whose total mass only represents 0.10 of the mass of all puffs will be subjected to enhanced merging.

Model changes to improve merging

A return code variable was added to the splitting routines to indicate to the main program that splitting has shutdown during the call to those routines (pufsph and pufspv). In the main program, each time splitting shuts down, FRHS is incremented by 0.5 to increase the effectiveness of puff merging, to a maximum value of 3 (FRHMAX in the namelist). Further, when splitting shuts down because the particle counter has approached the array limits, the puff age is changed to its negative value for those remaining puffs that are eligible to split but cannot due to the split restriction. Puffs with the split restriction are prevented from increasing in size (both horizontal and vertical) until the split restriction has been removed. This only occurs after the puffs split again (the puff number has been reduced by merging). Also at the first occurrence of the split restriction, the size to which a puff is permitted to grow before splitting is increased in proportion to the namelist parameter SPLITF. At the default value of 1.0, this value is automatically recomputed to be the ratio of the number of concentration grid cells to the maximum number of particles permitted. For values less than one this feature is disabled and splitting occurs as before, when the puff reaches the size of the meteorological grid (or concentration grid – whichever is larger). For values greater than one, that value is used to determine when a puff splits in proportion to the grid size. A termination message to standard output has been added prior to HYSPLIT completion if puff splitting restrictions are in place at the end of the simulation. Such a message would suggest that it might be necessary to rerun the simulation with added array space or different merge parameters.

Parameter Considerations

Very long duration simulations or simulations using very fine resolution meteorological data, which have an insufficient initial allocation of the puff array space (MAXPAR in the namelist) can result in split shutdown messages or perhaps even emission shutdown messages. If any of these occur, the simulation results should be viewed with caution. The results may be noisy and inaccurate if the emissions (new puffs released) have also been restricted. A simulation with puff split restrictions may be improved by first increasing the array space to a value that still results in acceptable simulation times. If not effective, or the CPU times become too long, the second choice could be to increase the frequency of enhanced merging (perhaps decreasing KRND from 6, to 3, 2, or even 1), and perhaps in combination with decreasing the split interval (increasing KSPLT from 1 to 3, 6, or 12). Although decreasing the split interval will not be effective once splitting has shutdown, it may extend the time at which splitting first shuts down. Enhanced merging has little effect if most of the puffs have the same mass, perhaps because they were released at the same time. It is most effective for continuous emission simulations, where there is a large range in puff mass due to the different number of splits each puff has been subjected. An effective removal method is setting FRMR to a non-zero value. However, this has the effect of purging the simulation of low-mass puffs and would be most appropriate for continuous emissions simulations, where the puffs at longer distances have less importance. Used incorrectly, setting this parameter to a non-zero value can seriously bias the model results. For long-duration continuous emission simulations, it may also be just as effective to stagger the emissions because it would not be necessary to emit puffs every time step for realistic (and accurate) results. This could be accomplished by emitting more mass over a shorter duration and then cycling the emissions. For instance instead of a continuous emission, one could emit 10 times the normal hourly amount over 0.1 hours (6 min) and the repeat the emission cycle (QCYCLE parameter) each hour. The emission cycle could even be staggered over longer times.

CPU Time Considerations

Long simulations may result in excessive CPU times because puffs will almost certainly be transported to the upper regions of the atmosphere where the winds are strongest which results in very small integration time steps. If computational accuracy in these upper regions is not required, perhaps because the only interest is boundary layer transport, the time step should be set to a fixed value. Given the same number of puffs, the enhanced merging version of the model should run substantially faster than the original version because when puff splitting shuts down and puffs continue to grow, they become quite large and cover many more concentration grid points which must all be sampled. Unrestricted puff splitting or restricting puff growth avoid this computational problem.

Theoretical Considerations

For example, if a global simulation is required using a 1-degree resolution concentration grid, that results in about 65,000 grid points at the surface. Clearly a very long duration simulation that is expected to spread over much of the domain will require a comparable number of puffs as grid cells to provide smooth concentration patterns. If we assume the lowest layer only represents about 10 percent of the volume over which the puffs have been transported and mixed, it is not unrealistic to expect such a simulation to require 10 times as many puffs. An alternate approach has been tested, although it is not yet part of the standard HYSPLIT distribution, would be to dump the puffs into a global grid model rather than splitting them as they grow to the size of the meteorological grid. Concentrations at any point would then be a sum from the two modeling approaches. This is a variation of a plume-in-grid model.

Test Simulations

All simulations were started at 45N 105W with a release of 50 particles over the first hour. The concentration grid resolution was one degree with 24 hour average outputs, the time step was fixed at 20 min, and the simulation duration was 480 hours using the 1-degree GDAS meteorological data.

Case 1: all versions, 3D particle, particles 500K/500K

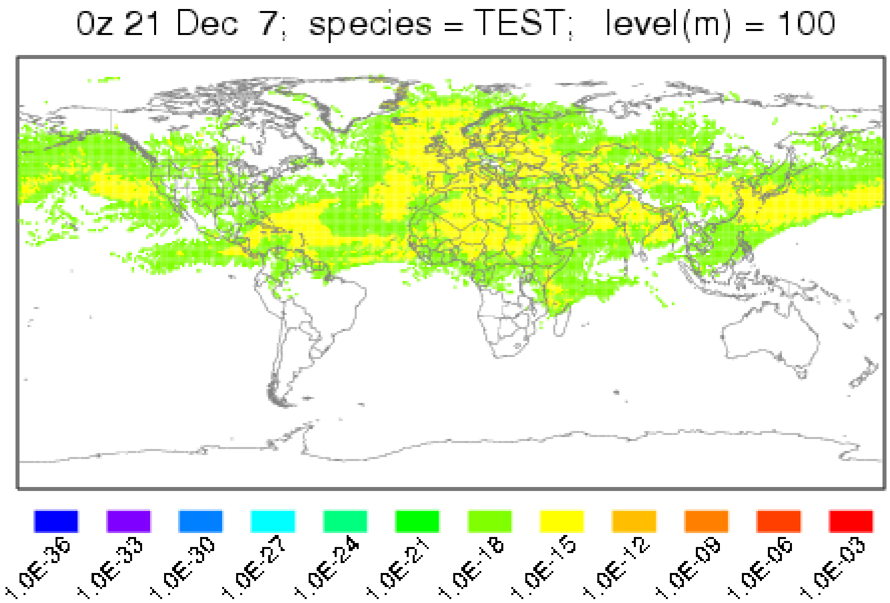

Case 1 represents the best possible simulation result. A sufficient number of particles were released the first hour (500,000) to properly cover the concentration domain after 20 simulations days. The calculation took 24 wall cock hours on a single ancient Power-4 processor.

Case 2: November 2007, default top-hat, particles 50/50K

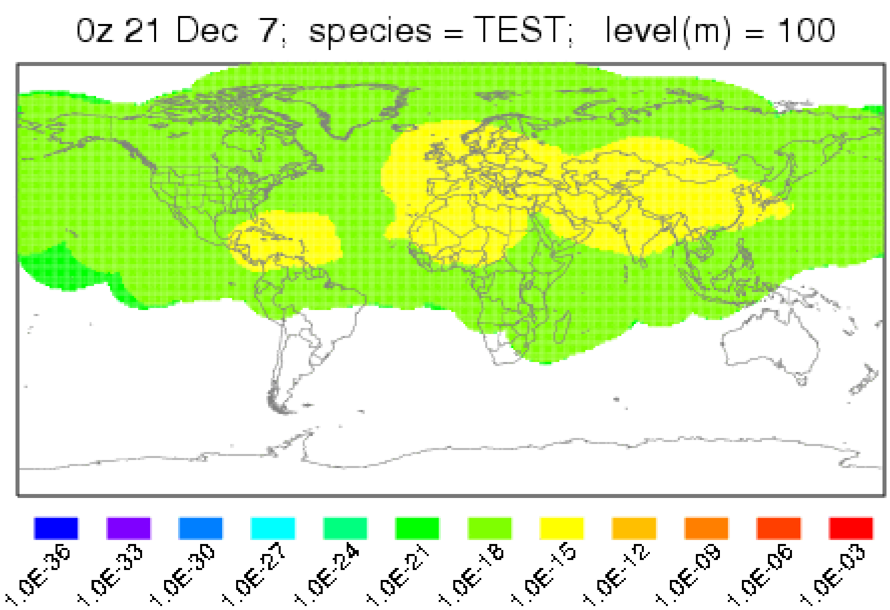

Case 2 shows the default results (top-hat particle) from the November 2007 and earlier versions of the model. Note the large smooth patterns, typical of when splitting shuts downs because the particle number has reached the array limits but puffs continue to grow with any restrictions.

Case 3: January 2008 update, auto enhanced merge, 50/50K

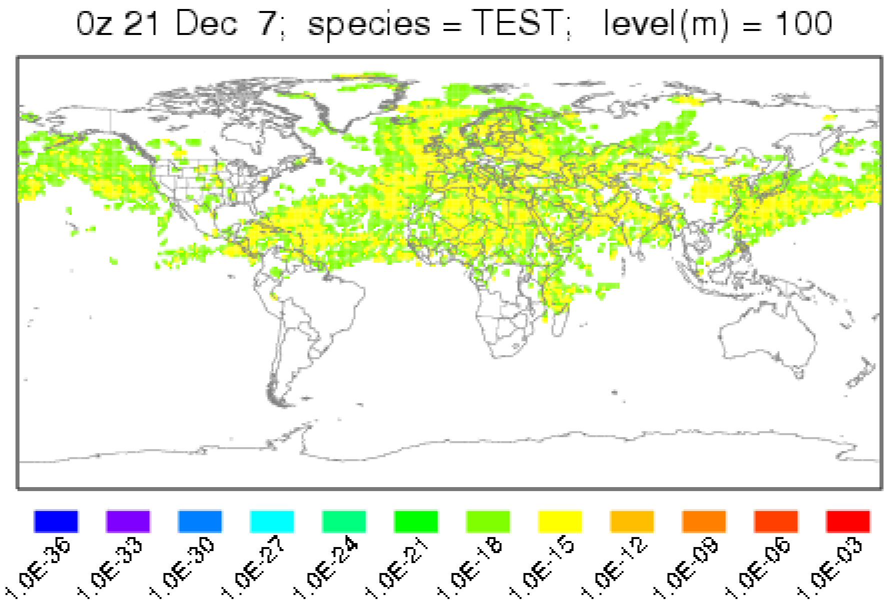

Case 3 shows the results from the January 2008 update with only the enhanced merging option turned on and with the puff growth restriction. Puffs still split when they reach the size of the meteorological grid. The concentration pattern is still very patchy compared with the 3d particle simulation.

Case 4: January 2008 update, delayed splitting with enhanced merge, 50/50K

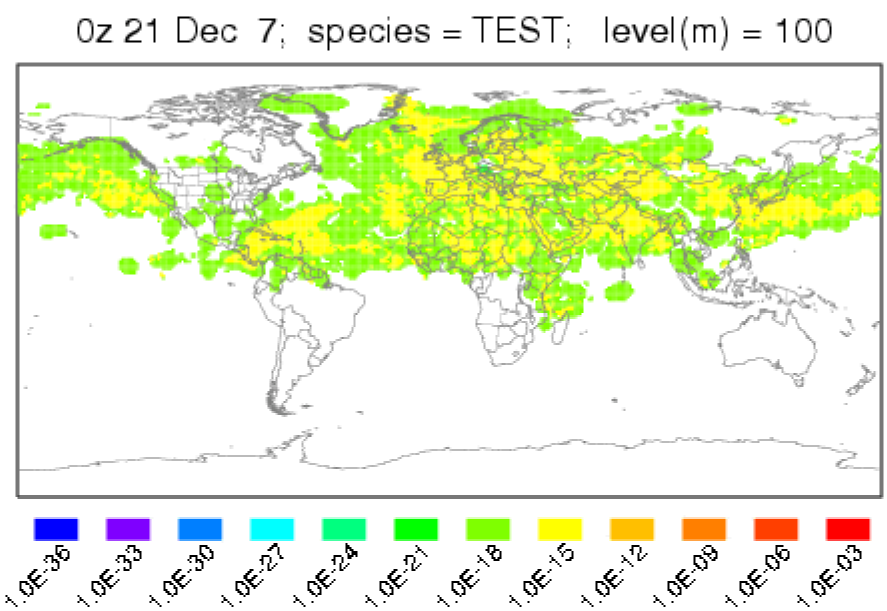

Case 4 represents the new standard calculation of the January 2008 update with enhanced merging turned on but delaying the puff splitting. Delayed puff splitting permits continued puff growth to provide a smoother concentration pattern without the need to turn off splitting. Smoother patterns can also be achieved by increasing the maximum particle number.