

5.2 Improve WRF I/O performance -- implementation of parallel IO.

Huang, Wei, *Hewlett Packard Enterprise Co.*

Many users know that WRF has quite few IO options: NetCDF 3/4, parallel-netcdf (pnetcdf), quilt, split, etc. i) NetCDF 3/4 works in general, but will be slow when domain is big in high resolution. ii) pnetcdf can speed up IO quite a bit, especially when MPI-IO is supported, and still have trouble with bigger domains. iii) Quilt may work for medium to large size domain, but it is tricky to setup, and need to take some processors out of total processors, which leaves less processor for computing. iv) Split is fast, but user has to run another program to stitch the split files together, and that could be slow. When there are many thousands of processors are used to run WRF, split generates many tons of small files, which could be a problem on some system, a simple "ls" command could take a very long time. WRF has another hidden IO options in 3.8.1, and 3.9.1, called Parallel IO (pio). Which can be manually configured to compile and run. We will present how pio works in WRF, and compare IO time with NetCDF3/4 and pnetcdf.