Using Domain Decomposition to Extend the DDSCAT Upper Limit of Target Size Parameter for Hydrometeors

Kwo-Sen Kuo*[(1)(2)], Robert W. Numrich[(3)], and Thomas L. Clune[(2)]
(1) Earth Systems Science Interdisciplinary Center, University of Maryland
College Park, Maryland, USA
(2) NASA Goddard Space Flight Center, Greenbelt, Maryland, USA
(3) The Graduate Center, City University of New York, New York, New York, USA

Fundamental to the remote sensing of atmospheric particulate matters, e.g. snow, rain, cloud ice, cloud water, and aerosol, is the single-scattering property of the particles involved. A significant portion of such particulate matter has complex shapes and therefore requires solution methods that provide more flexibility with regard to target shape than the Mie solution. The methods based on discrete dipole approximation, DDA, have proved to be particularly valuable in this regard, as they can accommodate practically any target shape. Among the publicly available DDA software, DDSCAT is considered to have an excellent balance between accuracy and performance.

The existing released versions of DDSCAT (7.2.x and prior), however, suffer from two important limitations. First, although they support execution on a distributed cluster through MPI, the number of parallel processes that can be used is limited to the number of independent orientations of the target. Second, their parallelization approach requires enormous amounts of memory, as each process is required to allocate memory for the full computational domain. In practice, large size parameters often require more memory than is available on a single node of a typical modern computing platform. This memory constraint can be partially ameliorated by using only one process per node, but only at the expense of idling all remaining CPU's and thereby significantly reducing the utilization efficiency of the computing resource.

We have constructed a variant of DDSCAT that directly addresses both of these limitations through the use of domain decomposition. In this new implementation, the MPI processes are subdivided into non-overlapping "teams" of processes, and each target orientation is solved by an entire team rather than a single process. The total number of processes is thus constrained by the product of the team size (number of processes, or team members, per team) and the number of orientations to be solved. Because the memory for the full computational domain is now distributed among the team members, large size parameters can be addressed simply by choosing suitably large teams. In essence, without our modifications, the maximum target size parameter that can be accommodated by DDSCAT is limited by the RAM available for each computing node; with them, it is limited by the RAM memory available for the entire cluster, i.e. sum of the RAM on all nodes. Furthermore, despite an increase in inter-process MPI communication, the application maintains strong scalability, i.e. the execution time for a fixed target size is approximately inversely proportional to the number of processors used. An additional improvement of this new implementation is a greater than 40% reduction in the execution time through the elimination of unnecessary FFT computations related to padding for the convolution.

We present some example DDSCAT results of very large target size parameters and further elaborate the implications of this new capability