Direct Domain Decomposition Methods (D^3M) for Electromagnetic Computations

Javad Moshfegh(1) and Marinos N. Vouvakis* (1)
(1) University of Massachusetts Amherst, Amherst, MA, 01003

Discretization of differential equations in electromagnetics produce very large but sparse linear systems that are often solved iteratively in lieu of memory demanding direct sparse matrix factorization methods. When the underlying models involve many independent excitations, or electrically large or complex or near resonance structures, even the most advanced preconditioned iterative methods lose effectiveness and efficiency. This problematic behavior, coupled with the recent proliferation of RAM memory has led to a recent resurgence of direct solution methods. Classical sparse Cholesky, LDL^T or LU factorization methods when coupled with advanced fill-in reducing reorderings such as multilevel nested dissection, and the multi-frontal or super-nodal cache aware computing algorithms can lead to very impressive performance. Most of these advanced algorithms are today available as “black-box” packages in the form of sparse matrix direct solver libraries, e.g. MUMPS, PARDISO, etc, and are heavily used by computational electromagnetics practitioners around the world.

This paper will outline an alternative exact direct solver paradigm that does not (entirely) rely on such “black-box” libraries, yet it can produce up to an order of magnitude less memory than those state-of-the-art solvers on problems discretized with unstructured tetrahedral tangential vector finite elements. More importantly, the proposed direct solution approach lends itself perfectly to parallel, distributed and heterogeneous CPU-GPU computing, paving the road for massive high-performance computing (HPC) direct solutions for FEM.

Drawing from our experiences in the area of iterative domain decomposition methods and advances in integral equation (IE) methods we will develop a direct domain decomposition method (D^3M) that instead of leveraging sparse matrix techniques that strive to minimize fill-ins, relies on smaller but partially dense matrices that strive for maximal data locality. Early results on 3D unstructured tetrahedron meshes and arbitrary volumetric geometries suggest that the proposed approaches can significantly outperform state-of-the-art sparse direct solvers in terms of memory usage while maintaining competitive serial implementation run times. Results of parallel implementation runs on multicore and distributed platforms will also be presented.