

# Parallelization of the Reduced-Coupling Technique for a Method-of-Moments-Based Field Solver Used for Product-Level Wide Data-Bus Analysis

A. J. Hesford<sup>1</sup>, J. D. Morsey<sup>2</sup>, W. C. Chew<sup>1</sup>, A. Deutsch<sup>2</sup>, and H. H. Smith<sup>3</sup>

<sup>1</sup> Dept. Electrical and Computer Engineering, University of Illinois at Urbana-Champaign  
1406 West Green Street, Urbana, IL 61801

Phone: (217) 333-0488, Email: hesford@uiuc.edu

<sup>2</sup> IBM T. J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, NY 10598

<sup>3</sup> IBM System and Technology Group, 2455 South Road, Poughkeepsie, NY 12601

**Abstract:** A parallel LU decomposition algorithm is presented to take advantage of the sparse impedance matrix produced by the reduced-coupling method. This algorithm allows rapid simulation of very large chip and packaging problems. A representative example is shown for a wide, on-chip data-bus that required one million surface unknowns and the computational power of a 1024-node IBM BlueGene cluster with distributed memory.

## Introduction

The reduced-coupling method was introduced in [1] to reduce the density of the Method of Moments (MoM) matrix. This is accomplished by subdividing the simulation structure into distinct regions, called groups. The basis functions wholly contained within a particular region are dubbed *internal* elements. Bases that overlap multiple regions are referred to as *boundary* elements. Boundary elements are restricted to overlap at most two regions, to avoid an overly-complicated formulation.

Full-physics interactions are represented between all internal basis functions within the same group. Likewise, each boundary element interacts with all of the internal elements in both groups sharing the boundary. Finally, a boundary element interacts with all other boundary elements that share a group with that element. Hence, the reduced-coupling impedance matrix may be represented as follows:

$$\bar{\mathbf{Z}}_{RC} = \begin{bmatrix} \bar{\mathbf{Z}}_{ii} & \bar{\mathbf{Z}}_{ib} \\ \bar{\mathbf{Z}}_{bi} & \bar{\mathbf{Z}}_{bb} \end{bmatrix}. \quad (1)$$

The sub-matrix  $\bar{\mathbf{Z}}_{ii}$  represents interactions between elements in the same group. With a proper indexing of the internal basis functions, this matrix will be block-diagonal. The matrices  $\bar{\mathbf{Z}}_{ib}$  and  $\bar{\mathbf{Z}}_{bi}$  represent interactions between boundary elements and internal elements. Because each internal element only interacts with boundary elements on the boundary of its own group, these matrices are both sparse. The matrix  $\bar{\mathbf{Z}}_{bb}$  represents interactions between boundary elements. Since boundary elements each interact with a restricted subset of all other boundary elements,  $\bar{\mathbf{Z}}_{bb}$  will also be sparse.

Figure 1 shows an example reduced-coupling arrangement. In this example, a total of five reduced-coupling blocks interact via six rows of boundary elements. In actual situations, each row of boundary elements may contain an arbitrary number of constituent elements. This example highlights the shape symmetry and sparsity of the  $\bar{\mathbf{Z}}_{ib}$  and  $\bar{\mathbf{Z}}_{bi}$  matrices.


Once  $\bar{\mathbf{Z}}_{RC}$  is constructed, it is desired to solve the general matrix equation

$$\begin{bmatrix} \bar{\mathbf{Z}}_{ii} & \bar{\mathbf{Z}}_{ib} \\ \bar{\mathbf{Z}}_{bi} & \bar{\mathbf{Z}}_{bb} \end{bmatrix} \cdot \begin{bmatrix} \bar{\mathbf{I}}_i \\ \bar{\mathbf{I}}_b \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{V}}_i \\ \bar{\mathbf{V}}_b \end{bmatrix}. \quad (2)$$

A solution to this equation is expressible in terms of the Schur complement of the block  $\bar{\mathbf{Z}}_{ii}$ :

$$\bar{\mathbf{I}}_b = \bar{\mathbf{Z}}_s^{-1} \cdot \left[ \bar{\mathbf{V}}_b - \bar{\mathbf{Z}}_{bi} \cdot \bar{\mathbf{Z}}_{ii}^{-1} \cdot \bar{\mathbf{V}}_i \right], \quad (3a)$$

$$\bar{\mathbf{I}}_i = \bar{\mathbf{Z}}_{ii}^{-1} \cdot \left[ \bar{\mathbf{V}}_i - \bar{\mathbf{Z}}_{ib} \cdot \bar{\mathbf{I}}_b \right], \quad (3b)$$



**Figure 1:** Sample reduced-coupling arrangement, with five groups interacting via six boundary elements. The sparsity of  $\bar{\mathbf{Z}}_{bb}$  is not shown.

where  $\bar{\mathbf{Z}}_s = \bar{\mathbf{Z}}_{bb} - \bar{\mathbf{Z}}_{bi} \cdot \bar{\mathbf{Z}}_{ii}^{-1} \cdot \bar{\mathbf{Z}}_{ib}$  is the Schur complement. When attempting to approximate the actual solution  $\mathbf{I}$  by the solution  $[\mathbf{I}_i, \mathbf{I}_b]^T$ , only the inverses of  $\bar{\mathbf{Z}}_{ii}$  and  $\bar{\mathbf{Z}}_s$  are required.

The reduced-coupling method was suggested in [2] as a preconditioner to a full-physics iterative solver. The promise of such a technique is a solution capturing the full physical behavior present in the impedance matrix  $\bar{\mathbf{Z}}$ , while hopefully reducing the number of iterations required for convergence. The sparsity of the reduced-coupling matrix makes it attractive for rapid inversion and application as a preconditioner.

## A Parallel LU Decomposition

Using the Schur complement to invert the preconditioner system still requires the inversion of  $\bar{\mathbf{Z}}_{ii}$  and its Schur complement,  $\bar{\mathbf{Z}}_s$ . A natural way to deal with these inverses is to compute the LU decomposition of each matrix, and use back-substitution when the inverse is needed. Because  $\bar{\mathbf{Z}}_{ii}$  is block-diagonal, so too is its LU decomposition. The decomposition of each block may proceed independently of other blocks, affording a low-cost parallel implementation.

However, it is often the case that each block of  $\bar{\mathbf{Z}}_{ii}$  is large. Furthermore, a large number of reduced-coupling groups often translates into a large number of boundary elements, resulting in a large  $\bar{\mathbf{Z}}_s$ . To further accelerate the parallel algorithm, the dense blocks of  $\bar{\mathbf{Z}}_{ii}$  are distributed amongst independent groups of processes. Likewise, the Schur complement  $\bar{\mathbf{Z}}_s$  is distributed amongst some subset of the processes assigned to the dense blocks.


Once the parallel LU decompositions of  $\bar{\mathbf{Z}}_{ii}$  and  $\bar{\mathbf{Z}}_s$  are complete, they may be used to solve (3). This is done with parallel forward- and back-substitution routines. The computational effort required to compute these solutions is negligible compared to that required by the LU decomposition.

## Modeling and Simulation of Wide Data-Bus on a Large Microprocessor Chip

It has previously been shown [5] that, due to the very fast transitions and short cycle times that are desired in high-performance microprocessor chips, there is a need for increased accuracy in timing and noise prediction. It is also explained in [5] that the finite impedance of the power distribution mesh for both voltage rails and signal lines has frequency-dependent properties and needs to be modeled accurately to capture the nonlinear interaction between simultaneous-switching noise, common-mode noise and crosstalk. Simplified techniques are being used in typical CAD tools due to the large problem size. In addition, it is important to include

such models in nonlinear simulations with actual driver and receiver circuits to accurately capture effects such as the compensating effects of power and interconnect noise, impact on timing, needed decoupling, use of multiple voltage rails and islands, and the effect of the effective impedance of the package.

Communication between processor and data caches involves very wide data-buses that need to draw large surges of current from the power distribution system. A methodology was described in [5] to partition the problem and utilize power-blocks with per-unit-length properties together with the interconnect transmission lines. The example that is described here shows, for the first time, the ability to perform a complete full-wave analysis of both the power distribution and the long global interconnects using the combined power of advanced accelerating field-solution techniques shown in [1, 3] and efficient parallelization on the fastest available supercomputer. The entire power mesh around a 12-line representative data-bus of lengths ranging from 0.25 to 2 mm, and including the topmost four wiring layers, was implemented as shown in Figure 2. The model required a discretization that resulted in between 175,000 to 1,050,000 surface unknowns. It was found in previous studies that such an analysis would need 50 times more unknowns when using finite-element type field solutions with volume discretization, or close to 50,000,000 unknowns.




**Figure 2:** Top-view representation of a 0.5 mm-length on-chip data-bus, with orthogonal wiring, used for full-wave electromagnetic analysis. A close-up image provides a look at 12 data lines.

## Numerical Results

A 12 line, 0.5 mm on-chip bus with orthogonal wiring and power distribution, as described in the previous section, was chosen as a benchmark problem, and scalability tested on the IBM BlueGene (BG) supercomputer. The BG system is composed of thousands of compute nodes, each with 512 MB of dedicated memory, linked by a fast communication network. When meshed, the on-chip structure is described by approximately 56,000 rectangular and 115,000 triangular mesh elements, yielding 288,000 MoM surface unknowns. This problem was solved using the reduced-coupling solution technique with 70 reduced-coupling regions, resulting in the parallelizable matrix topology discussed previously. The choice of problem, and therefore the plotted result, is of significance since this benchmark is not a canonical or academic problem designed specifically to produce ideal scaling, but is rather a real industrial model used to predict actual on-chip Delta-I and interconnect noise interactions.

The IBM-developed, MoM-based electromagnetic solver, EMSurf [3, 4], was used to solve for the 36-port S-parameter matrix over the 10 MHz to 30 GHz range on the IBM BG supercomputer, using between 128 compute nodes and 1024 compute nodes. The problem is not easily solved on fewer than 128 nodes, since each BG compute node has only 512 MB of memory and the reduced-coupling matrix for this problem requires approximately 21 GB of memory if stored in perfect sparse storage in a non-distributed fashion. The EM solver, with additional storage overhead, further increases the memory demands. The strict memory limits of each BG compute node place unique demands on the parallel algorithms employed, especially when performing very large simulations. Care must be taken to avoid unnecessary allocation of storage, and all parts of the simulation must operate in a distributed fashion.



**Figure 3:** Per-frequency solution time of MoM solution of on-chip data-bus using the parallel reduced-coupling method. Theoretical times are shown for algorithms exhibiting 100% and 0% parallel efficiency.

Figure 3 shows the scalability of the EM solution by plotting per-frequency solution time in seconds as a function of the number of processors used. For reference, the plots also show the theoretical solution times for algorithms exhibiting 100% and 0% parallel efficiency. The worst-case parallel efficiency for this data set was 47%. Additional data and larger examples will be presented.

## Conclusion

A parallel formulation of the reduced-coupling method has been presented for up to one million surface unknowns on a 1024-node supercomputer. Using this formulation, a parallel LU decomposition has been developed, which takes advantage of the block-diagonal nature of the intra-group interaction matrix  $\bar{\mathbf{Z}}_{ii}$ . The sparse nature of the matrix, coupled with a parallel LU decomposition algorithm, has allowed for the first time the solution of very large problems with a MoM-based field solver. This resulted in a great reduction of computation time and the complete, accurate analysis of large portions of product-level problems for broadband (10 MHz to 30 GHz) system prediction.

## References

- [1] B. J. Rubin, A. Mechentel, and J. H. Magerlein, "Electrical modeling of extremely large packages," in *Elec. Compon. Tech. Conf. Proc.*, San Jose, CA, 1997, pp. 804–809.
- [2] J. D. Morsey, "A broadband, low storage preconditioning scheme based on reduced coupling for full-wave method of moments solvers," in *Applied Computational Electromagnetics Society*, April 2004.
- [3] J. D. Morsey, B. J. Rubin, L. Jiang, L. Shan, L. B. Eisenberg, D. Becker, and M. Arsneault, "The use of fast integral equations solvers for practical package and interconnect analysis," in *15th Topical Meeting on Electrical Performance and Packaging (EPEP)*, Scottsdale, AZ, October 2006, pp. 335–338.
- [4] (2007) IBM AlphaWorks web site. [Online]. Available: <http://www.alphaworks.ibm.com/tech/eip>
- [5] A. Deutsch, H. H. Smith, B. J. Rubin, B. L. Krauter, and G. V. Kopcsay, "New methodology for the combined simulation of Delta-I noise interaction with interconnect noise for wide, on-chip data buses using lossy transmission-line power-blocks," *Trans. Adv. Pack.*, vol. 29, pp. 11–20, 2006.