Performance of low-rank QR approximation of the finite element Biot-Savart law

D. White, B. Fasenfest

October 20, 2006

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Performance of low-rank QR approximation of the finite element Biot-Savart law

Daniel A. White   and Benjamin J. Fasenfest

Abstract—

In this paper we present a low-rank QR method for evaluating the discrete Biot-Savart law. Our goal is to develop an algorithm that is easily implemented on parallel computers. It is assumed that the known current density and the unknown magnetic field are both expressed in a finite element expansion, and we wish to compute the degrees-of-freedom (DOF) in the basis function expansion of the magnetic field. The matrix that maps the current DOF to the field DOF is full, but if the spatial domain is properly partitioned the matrix can be written as a block matrix, with blocks representing distant interactions being low rank and having a compressed QR representation. While an octree partitioning of the matrix may be ideal, for ease of parallel implementation we employ a partitioning based on number of processors. The rank of each block (i.e. the compression) is determined by the specific geometry and is computed dynamically. In this paper we provide the algorithmic details and present computational results for large-scale computations.

Index Terms—

Maxwell's equations, Biot-Savart law, eddy currents, electromagnetic diffusion.

I. INTRODUCTION

The computation of magnetic fields from a prescribed electric current is a common problem in magnetic design and analysis. One approach is to formulate the problem as a Partial Differential Equation (PDE) for the unknown field with the prescribed electric current as the source term. Regardless of the particular PDE formulation, e.g. a magnetic vector potential formulation or a mixed *B-H* formulation, a large volumetric mesh must be employed, and some boundary condition must be applied on the outer boundary of the mesh. In contrast to the PDE approach, the Biot-Savart law can be employed to directly compute the magnetic field due to the prescribed current [1]. The advantage of the Biot-Savart law approach is that a full volume mesh is not required, and no boundary conditions need be applied. The disadvantage of the Biot-Savart approach is the computational cost, if there are $O(N)$ magnetic field observation points and $O(M)$ current samples the cost is $O(N*M)$. In this paper we review a fast low-rank QR method for compressing the $M \times N$ Biot-Savart matrix. The approach is similar to low-rank QR methods developed for boundary element electrostatics [2] [3] and for low frequency electric field integral equations [4]. The key difference with our approach is that we are concerned with volumetric current densities and implementation on parallel computers.

II. FORMULATION

The law of Biot and Savart is given by

$$\vec{B}(x) = \nabla \times \vec{A} = \frac{\mu}{4\pi} \int_{\Omega'} \frac{\vec{J}(x') \times (x - x')}{|x - x'|^3} d^3 x'. \quad (1)$$

where $J(x')$ is the known current density at the source point x', and $B(x)$ is the desired magnetic flux density at the observation point x. We assume that we have a finite element representation for J over the volume Ω', and a finite element representation for B over a surface Γ,

$$\vec{J} = \sum_{i=1}^{N} \xi_i \vec{W}_i^2 \ , \vec{B} = \sum_{j=1}^{M} \beta_j \vec{W}_j^1, \quad (2)$$

where ξ_j and β_j are the ith degree-of-freedom (DOF), and $\vec{W}_i^2$ and $\vec{W}_j^1$ are vector basis functions. Inserting the basis function expansions (2) into (1) yields the discrete Biot-Savart law

$$\mathbf{M}\bar{\beta} = \mathbf{Z}\bar{\xi}, \quad (3)$$

where

$$\mathbf{Z}_{ij} = \int_{\Gamma} \int_{\Omega'} \frac{\mu}{4\pi} \frac{\vec{W}_i^2(x') \times (x - x') \cdot \vec{W}_j^1(x)}{|x - x'|^3} d\Omega' \, d\Gamma, \quad (4)$$

and

$$\mathbf{M}_{ij} = \int_{\Gamma} \vec{W}_i^1(x) \cdot \vec{W}_j^1(x) \, d\Gamma \quad (5)$$

and where $\bar{\xi}$ and $\bar{\beta}$ are the arrays of DOF. The matrix $\mathbf{M}$ is a "mass matrix" due to the fact that the basis functions are not orthogonal. The mass matrix is extremely sparse and the computational cost for forming and solving this matrix is negligible. In many applications the problem of determining the *B*-field can be posed in terms of the magnetic vector potential $A = \nabla \times B$ with

$$\vec{A}(x) = \frac{\mu}{4\pi} \int_{\Omega} \frac{\vec{J}(x')}{|x - x'|} d^3 x'. \quad (6)$$

Using a finite element representation for *A* yields another version of the discrete Biot-Savart law

$$\vec{J} = \sum_{i=1}^{N} \xi_i \vec{W}_i^2 \ , \vec{A} = \sum_{j=1}^{M} \alpha_j \vec{W}_j^1, \quad (7)$$

$$\mathbf{M}\bar{\alpha} = \mathbf{Y}\bar{\xi}, \quad (8)$$

where

$$\mathbf{Y}_{ij} = \int_{\Gamma} \int_{\Omega'} \frac{\mu}{4\pi} \frac{\vec{W}_i^2(x') \cdot \vec{W}_j^1(x)}{|x - x'|} d\Omega' \, d\Gamma. \quad (9)$$

We will refer to the $M \times N$ matrices $\mathbf{Z}$ and $\mathbf{Y}$ as Biot-Savart matrices. The computation of these matrices involves singular and near-singular integrals. The surface integration is performed using standard Gaussian quadrature points for each surface element. The volume integration uses an adaptive integration rule, which varies the order of Gaussian quadrature based on the distance between the source point x' and the observation point x. When the surface element containing x is a face of the volume element containing x', a highly accurate height-based singularity cancellation quadrature rule is used [5]. The matrices (4) and (9) are constructed using 2-form or "face elements" for the basis functions W^2 and 1-form or "edge elements" for the basis functions W^1, see [6] for details on the construction of the basis functions.

Our primary application for the discrete Biot-Savart law is providing boundary conditions for finite element solution of multi-conductor eddy current problems. In each conductor we solve the time-dependent vector diffusion equation using an edge element based A-ϕ finite element formulation [7]. Clearly the B-field in the air surrounding the conductors is critical. The finite element formulation requires that either $\hat{n} \times \vec{A}$ or $\hat{n} \times \vec{B}$ be specified on the conductor boundaries, corresponding to inhomogeneous Dirichlet or Neumann boundary conditions, respectively. Our approach for dealing with the B-field in the air surrounding the conductors is to use the discrete Biot-Savart law (3) or (8) as the boundary condition on each conducting surface.

III. PARALLEL IMPLEMENTATION

We assume that the volume Ω has been partitioned into K partitions, where K is the number of computational processors, with each partition having an equal number of volume elements. The volume elements are distributed via the partitioning. The surface Γ is also partitioned into K equally sized surface partitions. Note however that the surface elements are not distributed via the surface partitions, each processor can access the entire surface mesh. The Biot-Savart matrix is then decomposed into a $K \times K$ block matrix, with every block $Z^{pq}, p \in \{1:K\}, q \in \{1:K\}$ representing the interaction of surface partition Γ_p with volume Ω_q. The qth processor computes blocks $Z^{pq}, p = 1 : K$, i.e. a column of blocks. Note that the matrix is decomposed via a partitioning of elements, hence the matrices Z^{pq} are overlapping in DOF space. The specific partitioning algorithm used to partition the elements is not critical, in the examples below we employ a graph-based algorithm [8]. The key point is that if the partitions Γ_p and Ω_q are well-separated then the sub-matrix Z^{pq} will have a low-rank QR decomposition. The procedure for computing the low-rank QR decomposition is described below. We define "well-separated" as follows: the bounding spheres for the element partitions Γ_p and Ω_q are computed, if the bounding spheres do not intersect then the partitions are considered well-separated and a low-rank QR representation of Z^{pq} is computed. We employ a recursive procedure for computing Z^{pq} when partitions Γ_p and Ω_q are not well-separated. This results in a hierarchical representation for Z. If Γ_p and Ω_q are not well separated, Ω_q is divided into eight equally sized sub-partitions, Γ_p is divided into four

equally sized sub-partitions, and the "well-separated test" is applied to the sub-partitions Γ_{pi} and Ω_{qj}, $i = 1 : 4, j = 1 : 8$. A space-filling curve algorithm is used for creating the sub-partitions. The process is applied recursively, with a low-rank QR representation computed for well-separated sub-partitions. The recursion is halted when a volume sub-partition contains fewer than some number of elements, for example 64 elements. If at the lowest level of recursion the interaction is not well separated, it is simply represented by a dense matrix. This is illustrated in Figure 1.

No parallel communication is required in the construction of the hierarchical Biot-Savart matrix, each processor has the elements that it needs to perform the integrals. Each processor has the same amount of work hence the computation of is load balanced. Note, however, that in the low-rank QR approximation the rank k is computed dynamically, and the rank k depends upon the geometry. Hence the application of the hierarchical Biot-Savart matrix, i.e. the matrix-vector multiplication $\bar{\beta} = \mathbf{Z}\bar{\xi}$, may not be perfectly load balanced. Also note that the application of the hierarchical Biot-Savart matrix does require parallel communication. This communication is as follows: (1) each processor q does a gather operation to get the values of $\bar{\xi}$ that it needs, (2) each processor q loops over the sub-matrices $Z^{pq}, p = 1 : K$ and computes $\bar{\beta}_q = \mathbf{Z}^{pq}\bar{\xi}_q$, (3) each processor participates in a global reduction on $\bar{\beta}_q$.


Fig. 1. Hierarchical partitioning of the Biot-Savart matrix. The highest level of partitioning is based on the number of processors. Some of the the the interactions at any *Level l* will be full rank (black boxes), and these interactions are sub-partitioned by decomposing the corresponding sub-volume and sub-surface to create *Level l+1*.

IV. LOW-RANK QR DECOMPOSITION

When Γ_p and Ω_q are well separated the matrix Z^{pq} will have a low-rank representation

$$Z^{pq}_{m \times n} \approx Q_{m \times k} \times R_{k \times n}, \tag{10}$$

where k is the rank. We do not want to form the entire Z^{pq} and then compress it, rather we sample the matrix by picking s rows and columns of Z^{pq}, where s is some predetermined number based on an estimate of the rank. The procedure for picking the sampled rows and columns is ad-hoc, the procedure that we employ is described in [4]. The sampling procedure is solely linear algebra, it does not depend upon the particular Green's function, finite element basis functions, etc. For the ad-hoc sampling procedure to be effective we must have s greater

than the expected rank. The algorithm for computing $Q_{m \times k}$ and $R_{k \times n}$ is as follows:

1: Form the sampled column matrix $S^c_{m \times s}$ and the sampled row matrix $S^r_{s \times n}$.

2: Compute the rank-revealing QR decomposition $\tilde{Q}_{m \times s} \tilde{R}_{s \times s} = S^c_{m \times s}$ using LAPACK routines DGEQPF and DORGQR. The rank k is determined by the criteria $\tilde{R}_{kk} < thresh \cdot \tilde{R}_{11}$ where *thresh* is a threshold value. Keep only k columns of $\tilde{Q}$, denote this as $Q_{m \times k}$, and discard $\tilde{R}$.

3: Form a new matrix $\hat{Q}_{s \times k}$ by taking s rows of $Q_{m \times k}$, the exact same rows as used to construct S^r.

4: Compute the least-squares solution to $\hat{Q}_{s \times k} R_{k \times n} = S^r_{s \times n}$ using LAPACK routine DGELSS.

At this point we have the desired matrices $Q_{m \times k}$ and $R_{k \times n}$ which approximate $Z^{pq}_{m \times n}$. To perform a matrix-vector multiplication with the compressed matrix it is necessary to include the permutations due to the column and row sampling, $Z^{pq}_{m \times n} \approx P^c \times Q_{m \times k} \times R_{k \times n} \times P^r$, where P^r and P^c are permutation matrices. The quality of the approximation, and the amount of compression (the rank k), are determined by the value of *thresh* used in Step 2 above. Our approach, being based on highly tuned LAPACK routines, is efficient both in terms of FLOPS and memory usage. The complexity of a single QR decomposition is $O(m \cdot s) + O(s \cdot n)$, using a fixed value of s yields a linear complexity in m and n.

The two key parameters in the QR decomposition are the threshold used to determine the rank k in Step 2, and the number of sampled rows and columns s in step 1. In practice we have found $0.01 < thresh < 0.001$ to yield acceptable results, meaning the the error in the QR compression is less than the intrinsic error of the finite element computation. Of course this is application dependent. For the parameter s we use a table lookup, where the argument is the normalized distance between the volume and surface regions. This normalized distance is defined as the distance between the centroids of the two bounding spheres divided by the sum of the radii of the spheres, a value of $d = 1$ means that the two bounding spheres are just touching. A table of computed ranks vs distance, generated by running a dozen different problems, is shown in Figure 2.


Fig. 2. Rank vs distance. The Black curve is the maximum rank, which is used to determine s. The Gray curve is the minimum rank.

V. EXAMPLES

In these examples we compute a hierarchical low-rank QR approximation of the matrix defined by Equation (9). For the first example consider the geometry shown in Figure 4. This geometry consists of 19000 volume elements and is partitioned for 16 processors. Therefore the Biot-Savart matrix will be a 16×16 block matrix. Each block Z^{pq} has roughly 1200 rows and 4000 columns. Using a value of $thresh = 0.005$ gives the parallel rank map shown in Figure 3. The compression is significant, each 1200×4000 matrix is compressed to $Q_{1200 \times k} + R_{k \times 4000}$ where k is the value shown in Figure 3. Note that the black blocks represent near interactions and have full rank. These near interaction blocks were decomposed further as explained in Section III above. The total compression was $60 \times$ for this specific example.


Fig. 3. Rank map for the highest level 16 x 16 partitioning of the three coil problem. Black=full rank, Dark Gray = $20 < k < 30$, Light Gray = $10 < k < 21$, White = $k < 11$.

The second example is shown in Figure 5. The geometry consists of three conducting coils, the center coil is driven with an independent current source, and we wish to compute the eddy currents in the coils due to the B-field in the surrounding air. The problem consists of 20736 volume elements and was partitioned for 24 parallel processors, therefore the Biot-Savart matrix is a 24×24 block matrix. Each processor had 24 matrices to compute at the highest level, on average 19 of these corresponded to well-separated regions and were compressed with an average rank of 10. The remaining 5 full-rank matrices were further partitioned into $4 \cdot 8 = 32$ sub-matrices, and on average 29 of these corresponded to well-separated regions and were compressed with an average rank of 25. At the lowest level of the hierarchy, the near interactions were represented, on average, by dense matrices of dimension 335×439, there were a total of $24 \cdot 3 = 72$ of these. The total compression was $109 \times$. This compression represents both the memory savings and the reduction in CPU time required to apply the Biot-Savart interaction.

We performed a scalability study by refining the computational mesh. As the mesh is refined, more processors were used to keep a constant number of mesh elements per processor. The results are shown in Figure 6, with a $N log_2(N)$ performance

for the memory usage. The CPU time for applying the compressed Biot-Savart matrix at every time step is proportional to the memory usage. The CPU time for computing the compressed Biot-Savart matrix is amortized over many thousands of time steps and is not dominant. Note that in Figure 6 the data for the uncompressed case is theoretical, since the memory usage was prohibitive.


Fig. 6. Compression of the hierarchical low-rank QR algorithm. The y-axis is words, the x-axis is number of volume elements, on a log scale. The solid line is $50 \cdot N log_2(N)$ for comparison.

Fig. 4. Computational mesh for a linear induction motor partitioned for 16 parallel processors. This partitioning gave a compression o $60\times$.

Biot-Savart law. The advantage of the proposed algorithm, which uses the parallel partitioning of the volume elements as the first level of the hierarchy, is that no parallel communication is require in the QR computation. Ideally the computational cost of computing the far-interactions will be less then cost of computing the near interactions, resulting in an $O(N)$ algorithm since the cost of the near interactions is $O(N)$. But it can be shown [3] that the cost of the far interactions will grow as $O(Nlog_2(N))$, and this was confirmed by our experiments. In addition, for a given mesh the performance will degrade as the number of processors p is increased. Given a volume partition of n elements and a surface partition of m elements, the QR work is $n \cdot k + m \cdot k$. If each of these partitions is cut in two (by doubling the number of processors p) the QR work is $4(n/2 \cdot k + m/2 \cdot k)$, the work is increased by a factor of 2.

VII. ACKNOWLEDGMENT

Fig. 5. Computational mesh for an inductive coupling application partitioned for 24 parallel processors. This partitioning gave a compression of $109\times$.

VI. CONCLUSIONS

Hierarchical low-rank QR compression of far interactions significantly decreases the cost of applying a finite element

REFERENCES

[1] J. D. Jackson. *Classical Electrodynamics*. 1962.
[2] S. Kapur and D. Long. IES3: Efficient electrostatic and electromagnetic solution. *IEEE Comp.. Sci. Eng.*, 5(4):60–67, 1998.
[3] D. Gope and V. Jandhyala. PILOT: A fast algorithm for enhanced 3d parasitic capacitance extraction. *Micro. Opt, tech. Lett.*, 41(3):169–173, 2004.
[4] D. Gope and V. Jandhyala. Efficient solution of EFIE via low-rank compression of multilevel predetermined interactions. *IEEE Trans. Ant. Prop.*, 53(10):3324–3333, 2005.
[5] M. A. Khaya and D. R. Wilton. Numerical evaluation of singular and near-signular potential integrals. *IEEE Trans. Ant. Prop.*, 53(10):3180–3190, 2005.
[6] P. Castillo, J. Koning, R. Rieben, and D. White. A discrete differential forms framework for computational electromagnetics. *Computer Modeling in Engineering & Sciences*, 5(4):331–346, 2004.
[7] R. Rieben and D. White. Verification of high-order mixed FEM solution of transient magnetic diffusion problems. *IEEE Trans. Mag.*, October 2005. article in press.
[8] G. Karypis and V. Kumar. A parallel algorithm for multilevel graph partitioning and sparse matrix ordering. *J. Parallel Distr. Comp.*, 48(1):71–95, 1998.