Parallel Multigrid Acceleration for the Finite Element Gaussian Belief Propagation Algorithm

Yousef El-Kurdi, Warren J. Gross, Dennis Giannacopoulos
Department of Electrical and Computer Engineering, McGill University,
Montreal (QC), H3A 0E9, Canada
yousef.el-kurdi@mail.mcgill.ca, warren.gross@mcgill.ca, dennis.giannacopoulos@mcgill.ca

***Abstract*—We introduce a novel parallel multigrid algorithm, referred to as FEM-MGGaBP, to accelerate the convergence of the recently introduced Finite Element Gaussian Belief Propagation solver. The FEM-MGGaBP algorithm processes the FEM computation in a fully distributed and parallel manner, element-by-element, demonstrating potential for high parallel efficiency. Our results for both sequential as well as parallel message scheduling versions of FEM-MGGaBP demonstrate high convergence rates independent of the scale of discretization on the finest mesh.**

***Index Terms*—finite element methods, multigrid, Gaussian belief propagation**

I. Introduction

Gaussian Belief Propagation (GaBP) is a distributed message passing algorithm originally used to compute marginal distributions on graphical models [1]. GaBP was also introduced as an algebraic solver for linear systems of equations [2], [3] that demonstrated great potential for robust parallel hardware implementations for Finite Element Methods (FEM) [4], [5]. In addition, GaBP has been shown to outperform classical iterative solvers such as Gauss-Seidel and Jacobi [2]. However such algorithms, which are derived based on pairwise interconnect assumptions on the underlying graphical model, suffer mostly from lack of convergence when diagonal dominance properties are not met. To address these convergence shortcomings and to improve the parallel efficiency of GaBP the FEM-GaBP algorithm was recently introduced, which was derived directly from a FEM variational formulation [6]. The FEM-GaBP algorithm solves the FEM in parallel element-by-element without the need to assemble a global sparse matrix. FEM-GaBP can be shown empirically to reach high parallel efficiency as the scale of the FEM problem increases [6]. However, like most iterative solvers, the FEM-GaBP convergence rate tends to stall when executed on fine meshes.

In this work we address this issue by introducing the novel Multigrid FEM-GaBP (FEM-MGGaBP) algorithm. Our empirical results of FEM-MGGaBP demonstrate that the new algorithm achieves high performance, typical of multigrid, which is independent of the scale of the finest grid. The FEM-MGGaBP achieves this performance even when executed using a new scheme of parallel message scheduling, promising the same parallel efficiency of the original FEM-GaBP. To our knowledge, this is the first multigrid formulation for continuous domain GaBP algorithms that is derived directly from a variational formulation of FEM.

II. Background

Multigrid accelerated solvers are among the fastest known algorithms to obtain solutions to linear systems of equations resulting from the FEM. The performance of multigrid can in practice, be shown to be independent of the size of the finest discretization of the domain [7]. Multigrid techniques can vary widely by using different relaxation algorithms with different smoothing properties, the approach of generating the course grid operators A^H, and the types of the interpolation and restriction transfer operators.

To generate multiple scales of fine meshes, we employ a hierarchical mesh refinement scheme based on element splitting. As shown in Fig. 1(a) each triangle is split into four geometrically similar child triangles by inserting nodes at midpoints of the parent triangle edges. This hierarchical scheme provides important advantages for element-by-element parallel behaviour of FEM-GaBP. In addition by utilizing semi-irregular mesh hierarchy, more adaptation to arbitrary domains can be achieved than using regular meshes.

III. Multigrid FEM-GaBP Algorithm

The distinguishing feature of the FEM-GaBP algorithm is solving the FEM in parallel, element-by-element, without needing to assemble the large sparse operator A or performing any global algebraic operations such as sparse matrix-vector multiplications. These key advantages, which have important implications for parallel hardware implementations, are maintained by the new multigrid formulation for FEM-GaBP. The FEM-GaBP solves the FEM by passing locally computed messages on a graphical model comprised of variable nodes and factor nodes [6]. The variable nodes correspond to the unknowns u, while the factor nodes correspond to probability functionals representing each finite element. Unlike classical relaxation algorithms such as Gauss-Seidel and Jacobi that are used in typical multigrid schemes, the FEM-GaBP algorithm operates by communicating local messages that have no direct correspondence with the global transfer operations of multigrid such as the residual and the correction.

To that end we define a quantity referred to as the factor node belief b_a, for each finite element a. The belief b_a takes the form of a multivariate Gaussian probability distribution such as:

$$b_a \propto \exp\left[-\frac{1}{2}u_a^T W_a u_a + K_a^T u_a\right] \quad (1)$$

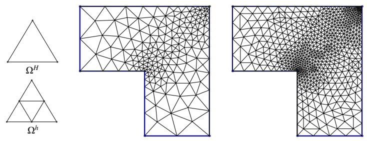

(a) (b) (c)

Figure 1. (a) Mesh refinement by splitting each triangle in mesh Ω^H into four geometrically similar sub-triangles to produce a finer mesh Ω^h. (b) Course irregular mesh. (c) Refined mesh by splitting.

where W_a is a small $n \times n$ dense matrix representing the inverse covariance of the factor node belief, K_a is a dense vector of dimension n, and u_a is a vector of random unknowns linked to the finite element a and having the same dimension. Matrix W_a and vector K_a are updated each iteration by incoming messages [6]. A key observation to state about the beliefs is that at message convergence, the joint mean vector of b_a, given by $\bar{u}_a = W_a^{-1} K_a$, will be equal to the marginal means of each of the random unknowns u_a computed by incorporating messages from all other connected factor nodes. Using this observation, we can formulate a quantity referred to as the belief residual. Using multigrids with hierarchical refinement by splitting, the belief residuals of each group of child triangles can be restricted into the parent triangle. This is a recursive and a local operation for each set of child-parent triangles. Also using the above belief definition, we can formulate a relationship between the coarse grid correction and the FEM-GaBP messages. Therefore, the resulting multigrid FEM-MGGaBP algorithm becomes a fixed point algorithm.

IV. Results

We illustrate the performance of FEM-MGGaBP by solving Laplace's equation for the L-shaped portion of the square coaxial-line problem shown in Fig. 1. The details of the problem will be provided in the long paper. A hierarchy of meshes is created by triangle splitting starting from an irregular course mesh. A V-cycle multigrid scheme is used where the parameters v_1 and v_2 are the number of pre-smoothing and post-smoothing iterations respectively. Since the FEM-MGGaBP operates on an element-by-element basis, the computational load is increased by a factor of four for each mesh refinement level. Table I shows a comparison of FEM-MGGaBP compared to the original FEM-GaBP algorithm. The solver is terminated when the relative error l^2-norm is dropped below 10^{-9}. The FEM-MGGaBP results demonstrate a multigrid acceleration performance that is independent of the number of unknowns on the finest level. This performance is illustrated by the amount of computational load reduction as the number of levels is increased. The computational reduction factors are computed by:

$$\text{Computational reduction} = \frac{\text{FEM-GaBP operations}}{\text{FEM-MGGaBP operations}}. \quad (2)$$

Fig. 2 shows the convergence rates for different pre-smoothing and post-smoothing settings as well as for different message scheduling schemes. The sequential message schedule provides the fastest convergence rates; however, it is not practical for parallel hardware implementations. It is important to note that the parallel scheduled FEM-MGGaBP has a comparable performance to the sequentially scheduled one showing that FEM-MGGaBP is sustaining high parallel efficiency. The parallel schedule scheme is based on a mesh triangle coloring approach, where messages of triangles of the same color are synchronized concurrently. This coloring scheme requires minimal overhead processing due to the utilized hierarchical mesh refinement scheme.

Table I

Computational reduction factors of the FEM-MGGaBP for the L-shaped conductor problem using four levels of refinement

Refinement Level	Num. variables	Tri-angles	FEM-GaBP relaxed	FEM-MGGaBP $v_1 = 1$, $v_2 = 1$	Computational reduction factors
1	825	1556	425	—	—
2	3205	6224	931	12	39
3	12633	24896	3663	12	122
4	50161	99584	13109	12	416

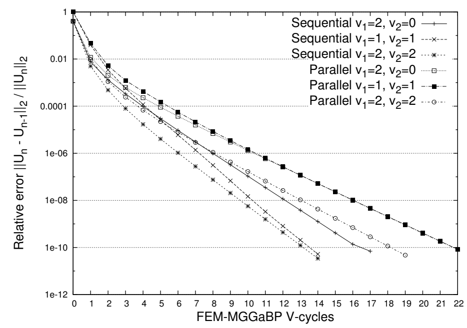

Figure 2. Sequential and parallel scheduled FEM-MGGaBP algorithm on four levels hierarchical mesh of the L-shaped conductor problem.

V. Conclusion

A novel FEM-MGGaBP algorithm was introduced and was shown to achieve high multigrid performance. The FEM-MGGaBP algorithm promises the same high parallel efficiency of the original FEM-GaBP algorithm. In the long version paper, we will provide a detailed formulation of the FEM-MGGaBP algorithm with more analysis on its performance. Also, we will detail the new parallel message scheduling algorithm which is based on a mesh coloring scheme.

References

[1] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.

[2] O. Shental, P. Siegel, J. Wolf, D. Bickson, and D. Dolev, "Gaussian belief propagation solver for systems of linear equations," in *IEEE Int. Symp. on Inform. Theory (ISIT)*, June–November 2008, pp. 1863–1867.

[3] Y. Weiss and W. T. Freeman, "Correctness of belief propagation in Gaussian graphical models of arbitrary topology," *Neural Computation*, vol. 13, no. 10, pp. 2173–2200, 2001.

[4] Y. El-Kurdi, W. Gross, and D. Giannacopoulos, "Efficient implementation of Gaussian belief propagation solver for large sparse diagonally dominant linear systems," *IEEE Trans. Magn.*, vol. 48, no. 2, pp. 471–474, 2012.

[5] Y. El-Kurdi, D. Giannacopoulos, and W. Gross, "Relaxed gaussian belief propagation," in *IEEE Int. Symp. on Inform. Theory (ISIT)*, July 2012, pp. 2002 –2006.

[6] Y. El-Kurdi, W. Gross, and D. Giannacopoulos, "Parallel solution of the finite element method using Gaussian belief propagation," in *The 15th Biennial IEEE Conference on Electromagnetic Field Computation*, 2012, p. 141.

[7] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A multigrid tutorial*, 2nd ed. Philadelphia, PA, USA: SIAM, 2000.