# Research Article 

# A Parallel Wavelet-Based Algebraic Multigrid Black-Box Solver and Preconditioner 

Fabio Henrique Pereira ${ }^{1}$ and Sílvio Ikuyo Nabeta ${ }^{2}$<br>${ }^{1}$ Industrial Engineering Post Graduation Program, Nove de Julho University (PMEP/UNINOVE), Francisco Matarazzo Avenue, 612, 05001100 São Paulo, SP, Brazil<br>${ }^{2}$ Electrical Machine and Drives Lab, São Paulo University (GMAcq/EP/USP), Luciano Gualberto Avenue, 380, 05508-010 São Paulo, SP, Brazil

Correspondence should be addressed to Fabio Henrique Pereira, fabiohp@uninove.br
Received 1 November 2011; Revised 26 January 2012; Accepted 8 February 2012
Academic Editor: Massimiliano Ferronato
Copyright © 2012 F. H. Pereira and S. I. Nabeta. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.


#### Abstract

This work introduces a new parallel wavelet-based algorithm for algebraic multigrid method (PWAMG) using a variation of the standard parallel implementation of discrete wavelet transforms. This new approach eliminates the grid coarsening process in traditional algebraic multigrid setup phase simplifying its implementation on distributed memory machines. The PWAMG method is used as a parallel black-box solver and as a preconditioner in some linear equations systems resulting from circuit simulations and 3D finite elements electromagnetic problems. The numerical results evaluate the efficiency of the new approach as a standalone solver and as preconditioner for the biconjugate gradient stabilized iterative method.


## 1. Introduction

The algebraic multigrid (AMG) method is one of the most efficient algorithms for solving large sparse linear systems. Especially in the context of large-scale problems and massively parallel computing, the most desirable property of AMG is its potential for algorithmic scalability: in the ideal case, for a matrix problem with $n$ unknowns, the number of iterative V-cycles required for convergence is independent of the problem size $n$ and the work in the setup phase and in each V-cycle is linearly proportional to the problem size $n$ [1, 2]. For all this, the need to solve linear systems arising from problems posed on extremely large, unstructured grids has been generating great interest in parallelizing AMG.

However, there are two major problems: first, the core of the AMG setup phase includes the grid coarsening process, which is inherently sequential in nature [1-3]. This coarsening scheme, for traditional AMG, can lead to computational complexity growth as
the problem size increases, resulting in an elevated memory use and execution time and in a reduced scalability $[4,5]$. Second, most parallel AMG algorithms are based on domain decomposition ideas, which have been proved to be very efficient but require a hierarchy of meshes that eliminates the algebraic characteristic of AMG and precludes its use as a blackbox method.

Due to those difficulties and the importance of the development of efficient parallel preconditioners for large, sparse systems of linear equations, the investigation of new parallel approaches has been the main subject of many researchers [6-12]. In this context, a great amount of work has been aimed to extract some parallelism from serial preconditioners such as factorization-based methods, which provide effective preconditioning on sequential architectures [8-10]. However, scalable parallel implementation of incomplete factorization preconditioners presents many limitations and challenges, and although some interesting approaches have presented good performance for certain classes of problems, quite scalable parallel algorithms for this kind of preconditioners seem to have not been available [8].

Also has received much attention in the last years the development of preconditioning approaches that have inherently parallel characteristics. In particular, approximate inverse preconditioners have proven extremely promising and effective for the solution of general sparse linear systems of equations [6, 7, 10-12]. Unfortunately, this kind of preconditioner also has some drawbacks: in general, as the discretization is refined the amount of work per grid point grows with problem size, and it is inherently difficult to approximate the inverse of very ill-conditioned linear systems with a sparse matrix [13].

In this work we introduce a new parallel algorithm for wavelet-based AMG (PWAMG) using a variation of the parallel implementation of discrete wavelet transforms. This new approach eliminates the grid coarsening process present at standard AMG setup phase, simplifying significantly the implementation on distributed memory machines and allowing the use of PWAMG as a parallel black-box solver and preconditioner. The parallel algorithm uses the message passing interface (MPI) that provides a standard for message passing for parallel computers and workstation clusters.

A sequential version of WAMG was introduced recently [14], and it has revealed to be a very efficient and promising method for several problems related to the computation of electromagnetic fields $[15,16]$. Here, the method is used as a parallel black-box solver and as a preconditioner in some linear equations systems resulting from circuit simulations and 3D finite elements electromagnetic problems. The numerical results evaluate the efficiency of the new approach as a standalone solver and as preconditioner for the biconjugate gradient stabilized iterative method.

## 2. The Discrete Wavelet Transform

The discrete wavelet transform (DWT) corresponds to the application of low-pass and high-pass filters, followed by the elimination of one out of two samples (decimation or subsampling). The discrete signal, which in one dimension is represented by a vector of values, is filtered by a set of digital filters that are associated to the wavelet adopted in the analysis.

Starting from a vector $y(N)$ at level 0 , two sets of coefficients are generated in each level $l$ of the process: a set $d_{l}$ of wavelets coefficients (detail coefficients) and a set $c_{l}$ of approximation coefficients. This procedure can be applied again, now using $c_{l}$ as an input vector to create new coefficients $c_{l+1}$ and $d_{l+1}$, successively.

Very important classes of filters are those of finite impulse response (FIR). The main characteristic of these filters is the convenient time-localization properties. These filters are originated from compact support wavelets, and they are calculated analytically. An example of FIR filters is the length-2 scaling filter with Haar or Daubechies-2 coefficients, which are given by (2.1):

$$
\begin{equation*}
h_{D 2}=\left[h_{0}, h_{1}\right]=\left[\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\right] . \tag{2.1}
\end{equation*}
$$

For more details about compact FIR filters see [17].
For a given vector $y=\left(y_{1}, y_{2}, \ldots, y_{N}\right)$, the Haar wavelet transform creates an approximation vector $c$ and a detail vector $d$ according to (2.2) and (2.3), respectively:

$$
\begin{align*}
& c=\left(c_{1}, c_{2}, \ldots, c_{N / 2}\right), \quad \text { with } c_{i}=\frac{\left(y_{2 i-1}+y_{2 i}\right)}{\sqrt{2}}  \tag{2.2}\\
& d=\left(d_{1}, d_{2}, \ldots, d_{N / 2}\right), \quad \text { with } d_{i}=\frac{\left(y_{2 i-1}-y_{2 i}\right)}{\sqrt{2}} \tag{2.3}
\end{align*}
$$

It is not difficult to see that this procedure will work only if $N$ is even.
In the 2D case, in which the discrete signal is represented by a matrix, the DWT is obtained through the application of successive steps of 1 D transform into the rows and columns of the matrix. This process generates a matrix formed by four types of coefficients: the approximation coefficients and the detail coefficients (horizontal, vertical, and diagonal), as illustrated in Figure 1. Blocks H and G represent, respectively, the low-pass and high-pass filters.

In both cases, the approximation coefficients keep the most important information of the discrete signal, whereas the detail coefficients possess very small values, next to zero. These approximation coefficients will contain low-pass information, which is essentially a low-resolution version of the signal and represent a coarse version of the original data.

## 3. A Wavelet-Based Algebraic Multigrid

The approximation property of wavelet is explored by wavelet-based algebraic multigrid for creating the hierarchy of matrices. The method considers the use of a modified discrete wavelet transform in the construction of the transfer operators and the hierarchy of matrices in the multigrid approach.

A two-dimensional modified DWT is applied to produce an approximation of the matrix in each level of the wavelets multiresolution decomposition process. An operator formed only by low-pass filters is created, which is applied to the rows and columns of the matrix. This same operator is used for the intergrid transfer in the AMG. If a length- 2 (first order) scaling filter is used, for example, which means the application of the operation


Figure 1: The two-dimensional DWT: one-dimensional transform in the rows and columns of the matrix.
defined in (2.2) in the rows and columns of the matrix, the operation is matricially defined as (3.1):

$$
P_{k}^{k+1}=\left(\begin{array}{ccccccccc}
h_{1} & h_{0} & 0 & 0 & 0 & \cdots & \cdots & \cdots & 0  \tag{3.1}\\
0 & 0 & h_{1} & h_{0} & 0 & 0 & 0 & \cdots & 0 \\
\vdots & & & & \vdots & & & & \vdots \\
0 & 0 & 0 & \cdots & \cdots & \cdots & 0 & h_{1} & h_{0}
\end{array}\right)_{(N / 2) \times N} .
$$

The prolongation operator is defined in the usual form,

$$
\begin{equation*}
P_{k+1}^{k}=\left(P_{k}^{k+1}\right)^{T}, \tag{3.2}
\end{equation*}
$$

and the matrix in the corresponding level $k$ with the Galerkin condition:

$$
\begin{equation*}
A^{k+1}=P_{k}^{k+1} A^{k} P_{k+1}^{k}, \tag{3.3}
\end{equation*}
$$

reminding that $A^{0}=A$ is the matrix of the original system.
Once the intergrid transfer operators and the hierarchy of matrices are created, the multigrid method (V-cycle) can be defined as usual, using a recursive call of the following two-level algorithm.

Algorithm 3.1 (Two-level multigrid). The following holds.
Input: the right hand side vector $b$, the original matrix $A$
and the coarse grid matrix $\mathrm{PAP}^{T}$
Output: approximation $\tilde{x}$
(1) Choose an initial guess $x$ and apply $\mathcal{\nu}_{1}$ smoothing steps in $A x=b$
(2) Compute $r=b-A x$
(3) $e=\left(\mathrm{PAP}^{T}\right)^{-1} \mathrm{Pr}$
(4) $\tilde{x}=x+P^{\mathrm{T}} e$
(5) Apply $v_{2}$ smoothing steps in $A \tilde{x}=b$
(6) Return $\tilde{x}$.

In this paper, a V-cycle multigrid with $\nu_{1}=v_{2}=1$ is applied iteratively as a solver and also as a preconditioner inside the biconjugate gradient stabilized (BiCGStab) iterations.

### 3.1. The Choice of the Wavelet Filters

A common problem on the choice of the filters is to decide between the fill-in control and the wavelet properties. As the WAMG often deals with sparse matrices, the control of the nonzero number is a very important task. In this case, if the matrix $A^{k}$ is sparse, then the number of nonzero elements in the next level matrix $A^{k+1}=P_{k}^{k+1} A^{k} P_{k+1}^{k}$ will depend on the order of the filter used in the restriction and prolongation operators. In fact, the longer the filter used the larger the number of nonzero entries in the next computed matrix. Consequently, most of the wavelet-based multigrid methods use shorter filters such as Haar or Daubechies-2 coefficients in its approaches [18-23]. This is also the case in this paper.

## 4. A Parallel Wavelet-Based Algebraic Multigrid

One advantage of the proposed approach is the elimination of the coarsening scheme from the setup phase. The application of the operations defined in (2.3) in the rows and columns of a matrix allows creating an approximation without any information about meshes or the adjacency graph of the matrix. Thus the implementation on distributed memory machines becomes simpler allowing the use of the method as a parallel black-box solver and preconditioner. Our approach, based on the parallel implementation of discrete wavelet transform presented in [24], avoids any communication among processors in the setup phase if first-order filters are used.

The parallelization strategy starts dividing equally the number $N$ of rows of the matrix between the processors, in such a way that the 1-D transform defined in (2.3) could be applied entirely locally in the rows. As each matrix column is equally divided, the transformation in this direction is accomplished partially in each processing unit, considering the size of the column in its local memory. For $n p$ processors, for example, each processor will apply the wavelet transform in the $\lceil N / n p\rceil$ elements of the column in its memory, where $\lceil x\rceil$ means the largest integer even less than $x$. If the largest integer less than $N / n p$ is odd, then the last element of the column in the local memory is unchanged, as illustrated in Figure 2 for $N=10$ and $n p=2$. In the figure, $c_{i, j}^{l}$ means the $i$ th local element of the level $l$ vector $c$ in the processor $j$.

Processor no. 1

Processor no. 2


Figure 2: Illustration of the 1D parallel wavelet transform with 2 processors.


Figure 3: Matrix-vector product of a matrix distributed by rows.


Figure 4: MPI collective communication function.

Thus the resulting coarse level matrix is calculated entirely locally.
In the solver phase the interprocessors communication is necessary only for operations involving matrices. More specifically, it is necessary to update the vector always after the smoothing step and after the matrix-vector product in the residual calculation (lines 1, 2, and 5 of the algorithm). In a matrix-vector product $A * b=c$, for example, the matrix $A$ is distributed in rows, the vector $b$ is shared by all processors, and vector $c$ is calculated in parallel as illustrated in Figure 3, for 4 processors. Then the resulting vector $c$ is updated by the processors through the message passing interface library (MPI). This task is accomplished by using the MPI collective communication function MPI Allgather [25]. The MPI Allgather function effect is shown in Figure 4. It is important to highlight that only the resulting vector should be updated. It means each processor communicates to the other only a few elements


Figure 5: The Beowulf Linux cluster architecture.

Table 1: Electromagnetic matrices properties.

| Matrix properties | 2cubes_sphere | Offshore | dielFilterV2real | Circuit5M_dc |
| :--- | :---: | :---: | :---: | :---: |
| Number of rows | 101,492 | 259,789 | $1,157,456$ | $3,523,317$ |
| Number of columns | 101,492 | 259,789 | $1,157,456$ | $3,523,317$ |
| Nonzeros | $1,647,264$ | $4,242,673$ | $48,538,952$ | $14,865,409$ |
| Explicit zero entries | 0 | 0 | 0 | $4,328,784$ |
| Type | Real | Real | Real | Real |
| Structure | Symmetric | Symmetric | Symmetric | Unsymmetric |
| Positive definite? | Yes | Yes | No | No |

Table 2: Results for sequential tests (not enough memory).

| Matrix | ILU + BiCGStab |  | WAMG |  |  | WAMG + BiCGStab |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Setup | Solver | $n$ | Setup | Solver | $n$ | Setup | Solver | $n$ |
| Circuit5M_dc | 6.96 | 29.98 | 3 | 47.4 | 14.37 | 3 | 61.30 | 45.49 | 2 |
| 2cubes_sphere | 1.98 | 2.15 | 3 | 5.29 | 2.02 | 3 | 5.30 | 1.55 | 4 |
| Offshore | 5.36 | 21.55 | 10 | 16.67 | 18.42 | 6 | 16.69 | 21.17 | 10 |
| dielFilter-V2real | - | - | - | - | - | - | - | - | - |

that it stores. However, in order for the process to continue, the whole vector must be updated on each processor and some kind of collective communication should take place.

The PWAMG method in this work uses a hybrid Jacobi-Gauss method as smoother and the V-cycle for the resolution scheme. The Gauss-Seidel smoothing is applied inside each processor and the Jacobi method applied interprocessors.

## 5. The Numerical Test Problems

The parallel algorithm uses the version one of the MPI that provides a standard for message passing for parallel computers and workstation clusters. The method has been implemented

Table 3: Parallel results for 2cubes_sphere matrix.

using C++ and tested in a homogeneous Beowulf cluster with 6 machine nodes (Core 2 Duo, 1 GB RAM) connected to a switched fast Ethernet network, as illustrated in Figure 5.

The multigrid V-cycle approach was applied as a resolution scheme to solve some linear equations systems resulting from circuit simulations and finite elements electromagnetic problems. The parallel method was also applied as a preconditioner for the iterative biconjugate gradient stabilized (BiCGStab) method, which has been implemented using the same vector updating approach.

The three first matrices are related to 3D finite element electromagnetic analysis. The matrices 2 cubes sphere and offshore are from a study of edge-based finite-element time domain solvers for 3D electromagnetic diffusion equations. The dielFilterV2real is a real symmetric matrix, which comes from a high-order vector finite element analysis of a 4th-pole dielectric resonator [26]. The last real unsymmetric matrix is from circuit simulation problems from Tim Davis sparse matrix collection [27]. The matrices properties are presented in Table 1.

## 6. Results

The problems were solved firstly using the sequential version of the proposed method. For comparison, the BiCGStab method preconditioned by the incomplete LU was used.

Table 4: Parallel results for offshore matrix.


The results are presented in Table 2. In all cases, the convergence is defined by $\left\|r^{n}\right\| /\|b\|<10^{-6}$, where $r^{n}$ is the residual vector at the $n$th iteration and the right hand side vector $b$ is chosen so that the solution is a unitary vector. The setup and solver times are in seconds, and $n$ is the number of iterations.

The parallel results for the circuit simulation and electromagnetics problems are presented in Tables 3, 4, 5, and 6 for the PWAMG solver and preconditioner. The setup and solver times are in seconds and have been presented in the form $t 1$ ( $t 2$ ), where $t 1$ is the processing time returned by the C clock() function and $t 2$ is the total time spent in the corresponding phase, which includes the MPI communication time and is measured using the MPI function MPI_Wtime().

In some cases the CPU and wall-clock times are nearly equal, which means that the processes that are waiting for synchronization are still consuming full CPU time. Moreover, the CPU time and the wall-clock time are practically the same in setup phase indicating that there are no interprocessor communications during the setup. It is a great advantage of the proposed approach since this phase is the most time-consuming task in multilevel approaches, as can be seen in Table 2. For single processor algorithms the times $t 1$ and $t 2$ are the same.

Table 5: Parallel results for circuit5M_dc matrix.


Three important aspects related to the parallel results should be highlighted.
(1) The proposed method uses a hybrid Jacobi Gauss-Seidel method as a smoother and a coarse system solver. This method applies the Gauss-Seidel method inside each processor and the Jacobi method between processors. So, as the number of processors increases, both smoother and coarse solvers become different. Also in the sequential version of the method this approach has not been reproduced. This can help us to explain the difference in the number of iterations in the parallel and sequential versions.
(2) The way in which the matrices are split between the processors (the same number of rows) may cause load unbalance in some cases that can affect the overall performance of the method. For the matrix circuit5M_dc, for example, the number of nonzero elements in the processor one is $25 \%$ larger than in the processor two, when 2 processors are used. A new way to divide the matrix among the processors should be investigated.
(3) Despite the use of a homogeneous Beowulf cluster, the type of network connection based on a switched fast Ethernet network shows to be an important bottleneck. So, the effects of the fast Ethernet connection on the results should be evaluated separately. In order to do so, the time $(p)$ required by the solver phase execution on $p$ processors, in which there is interprocessors communication, was evaluated
considering a relative time in relation to the MPI $\pi$ calculation example [25]. It means the time $(p)$ is defined as (6.1)

$$
\begin{equation*}
\operatorname{time}(p)=\frac{w \operatorname{time}(p)}{\pi-\operatorname{time}(p)} \tag{6.1}
\end{equation*}
$$

in which wtime $(p)$ is the total time spent in the setup phase, which includes the MPI communication time and is measured using the MPI function MPI_Wtime(), and $\pi$-time $(p)$ is the time spent by the MPI $\pi$ example, both with $p$ processors. As an example, Table 7 presents the values of $\pi$-time $(p)$, wtime $(p)$, $\operatorname{time}(p)$ and $\operatorname{speedup}(p), p=1, \ldots, 6$, for the matrix 2cubes_sphere, which were used to create the illustration in Figure 6(a). The values of $\pi$-time $(p)$ were obtained as the mean of 5 runs. All the other results in Figure 6 were derived in a similar way.

As usual, the absolute speedup $S_{p}$ is used for analyzing the parallel performance, and it is defined as (6.2)

$$
\begin{equation*}
S_{p}=\operatorname{speedup}(p)=\frac{\operatorname{time}(1)}{\operatorname{time}(p)} \tag{6.2}
\end{equation*}
$$

in which time (1) is the time spent by the best sequential algorithm and time $(p)$ as defined in (6.1).

As the MPI $\pi$ example uses only point-to-point communication functions, the relative time approach can help to clarify if the apparent poor scalability is due to a high collective communication cost or mainly due to the type of network connection used.

The solver and preconditioner speedups are illustrated in Figure 6 for the matrices circuit5M_dc, 2cubes_sphere, and offshore.

## 7. Conclusions

The PWAMG method, proposed in this work, has been applied as a black-box solver and preconditioner in some circuit simulations and finite element electromagnetic problems with good results.

An important characteristic of the proposed approach is its small demand for interprocessors communication. Actually, no communication is required in the setup phase if first-order filters are used. This characteristic is confirmed by the results for setup time presented in Tables 3-6, observing that the times measured by MPI_Wtime() and C clock() functions are practically the same. It is a great advantage of the proposed approach since this phase is the most time-consuming one.

In the solver phase, in which there is interprocessors communication, the numerical results seem to show a poor performance with more than 2 processors, even when the number of iterations does not increase by using more processors. These results can be caused in part due to the use of a collective communication function to update the vector after the matrix operations, which is motivated by the manner the matrix is divided between the processors.

In order to update the vector each processor should send and receive the part of the vector to all of the other processors. Of course, when the number of processors increases, this work also becomes larger.

Table 6: Parallel results for dielFilterV2real matrix (not enough memory).


Table 7: Solver phase speedup using a relative time (p): example for matrix 2cubes_sphere.

|  | $p=1$ | $p=2$ | $p=3$ | $p=4$ | $p=5$ | $p=6$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| (A) $\operatorname{\text {wime}}(p)$ | 2.02 | 2.99 | 3.47 | 3.26 | 3.32 | 3.27 |
| (B) $\pi$-time $(p)$ | 0.0000885 | 0.0002926 | 0.0003222 | 0.000997 | 0.001011 | 0.0009938 |
| (C) $\operatorname{time}(p)=$ A $/ \mathrm{B}$ | 22824.86 | 10218.73 | 10769.71 | 3269.81 | 3283.88 | 3290.40 |
| (D) speedup $(p)$ | 1.00 | 2.23 | 2.12 | 6.98 | 6.95 | 6.94 |

An alternative to overcome this problem may be to apply some graph partitioning method and develop an approach in which each processor should communicate only with its neighbors. Such approach is under development, and it is out of the scope of this paper.

However, the type of network connection based on a switched fast Ethernet network shows to be an important bottleneck, and its effects should be evaluated separately. As the MPI $\pi$ example uses only point-to-point communication functions, the relative time approach can help to clarify if the apparent poor scalability is due to a high collective communication cost or mainly due to the type of network connection used. The speedup results based on the relative times show that the proposed approach is promising and that the kind of network connection that has been used is maybe the most important drawback.

Moreover, in spite of the speedup being less than the linear in some cases, it is important to mention an important aspect of this application: in the context of large sparse linear system of equations, where this paper is inserted, the problems have large memory


Figure 6: Speedup for 2cubes_sphere solver (a) and preconditioner (b), offshore solver (c) and preconditioner (d), and circuit5M_dc solver (e) and preconditioner (f).
requirements. In these cases, as presented in [28], the speedup necessary to be cost effective can be much less than linear. The parallel program does not need $p$ times memory of the unit processor, since parallelizing a job rarely multiplies its memory requirements by $p$.

Finally, the number of iteration of the proposed method seems to be independent of the number of processors. However, it is necessary to carry out more tests with a larger number of processors in order to draw more definitive conclusions. Nowadays, the authors are looking for new resources and/or partnerships to enable the continuation of this work.

## Acknowledgments

Authors would like to thank the São Paulo Research Foundation (FAPESP) (06/59547-5 and 2011/06725-1) and Nove de Julho University (UNINOVE) for the financial support.

## References

[1] G. Haase, M. Kuhn, and S. Reitzinger, "Parallel algebraic multigrid methods on distributed memory computers," SIAM Journal on Scientific Computing, vol. 24, no. 2, pp. 410-427, 2002.
[2] V. E. Henson and U. M. Yang, "BoomerAMG: a parallel algebraic multigrid solver and preconditioner," Applied Numerical Mathematics, vol. 41, no. 1, pp. 155-177, 2002.
[3] H. de Sterck, U. M. Yang, and J. J. Heys, "Reducing complexity in parallel algebraic multigrid preconditioners," SIAM Journal on Matrix Analysis and Applications, vol. 27, no. 4, pp. 1019-1039, 2006.
[4] M. Griebel, B. Metsch, D. Oeltz, and M. A. Schweitzer, "Coarse grid classification: a parallel coarsening scheme for algebraic multigrid methods," Numerical Linear Algebra with Applications, vol. 13, no. 2-3, pp. 193-214, 2006.
[5] A. J. Cleary, R. D. Falgout, V. E. Henson, and J. E. Jones, Coarse-Grid Selection for Parallel Algebraic Multigrid, Lawrence Livermore National Laboratory, Livermore, Calif, USA, 2000.
[6] L. Yu. Kolotilina and A. Yu. Yeremin, "Factorized sparse approximate inverse preconditionings. I. Theory," SIAM Journal on Matrix Analysis and Applications, vol. 14, no. 1, pp. 45-58, 1993.
[7] M. J. Grote and T. Huckle, "Parallel preconditioning with sparse approximate inverses," SIAM Journal on Scientific Computing, vol. 18, no. 3, pp. 838-853, 1997.
[8] D. Hysom and A. Pothen, "A scalable parallel algorithm for incomplete factor preconditioning," SIAM Journal on Scientific Computing, vol. 22, no. 6, pp. 2194-2215, 2000.
[9] P. Raghavan, K. Teranishi, and E. G. Ng, "A latency tolerant hybrid sparse solver using incomplete Cholesky factorization," Numerical Linear Algebra with Applications, vol. 10, no. 5-6, pp. 541-560, 2003.
[10] P. Raghavan and K. Teranishi, "Parallel hybrid preconditioning: incomplete factorization with selective sparse approximate inversion," SIAM Journal on Scientific Computing, vol. 32, no. 3, pp. 13231345, 2010.
[11] C. Janna, M. Ferronato, and G. Gambolati, "A block Fsai-Ilu parallel preconditioner for symmetric positive definite linear systems," SIAM Journal on Scientific Computing, vol. 32, no. 5, pp. 2468-2484, 2010.
[12] C. Janna and M. Ferronato, "Adaptive pattern research for block FSAI preconditioning," SIAM Journal on Scientific Computing, vol. 33, no. 6, pp. 3357-3380, 2011.
[13] M. Benzi and M. Tuma, "A comparative study of sparse approximate inverse preconditioners," Applied Numerical Mathematics, vol. 30, no. 2-3, pp. 305-340, 1999.
[14] F. H. Pereira, S. L. L. Verardi, and S. I. Nabeta, "A wavelet-based algebraic multigrid preconditioner for sparse linear systems," Applied Mathematics and Computation, vol. 182, no. 2, pp. 1098-1107, 2006.
[15] F. H. Pereira, M. F. Palin, S. L. L. Verardi, V. C. Silva, J. R. Cardoso, and S. I. Nabeta, "A wavelet-based algebraic multigrid preconditioning for iterative solvers in finite-element analysis," IEEE Transactions on Magnetics, vol. 43, no. 4, pp. 1553-1556, 2007.
[16] F. H. Pereira, M. M. Afonso, J. A. De Vasconcelos, and S. I. Nabeta, "An efficient two-level preconditioner based on lifting for FEM-BEM equations," Journal of Microwaves and Optoelectronics, vol. 9, no. 2, pp. 78-88, 2010.
[17] T. K. Sarkar, M. Salazar-Palma, and C. W. Michael, Wavelet Applications in Engineering Electromagnetics, Artech House, Boston, Mass, USA, 2002.
[18] V. M. Garcıa, L. Acevedo, and A. M. Vidal, "Variants of algebraic wavelet-based multigrid methods: application to shifted linear systems," Applied Mathematics and Computation, vol. 202, no. 1, pp. 287299, 2008.
[19] A. Avudainayagam and C. Vani, "Wavelet based multigrid methods for linear and nonlinear elliptic partial differential equations," Applied Mathematics and Computation, vol. 148, no. 2, pp. 307-320, 2004.
[20] D. de Leon, "Wavelet techniques applied to multigrid methods," CAM Report 00-42, Department of Mathematics UCLA, Los Angeles, Calif, USA, 2000.
[21] G. Wang, R. W. Dutton, and J. Hou, "A fast wavelet multigrid algorithm for solution of electromagnetic integral equations," Microwave and Optical Technology Letters, vol. 24, no. 2, pp. 86-91, 2000.
[22] R. S. Chen, D. G. Fang, K. F. Tsang, and E. K. N. Yung, "Analysis of millimeter wave scattering by an electrically large metallic grating using wavelet-based algebratic multigrid preconditioned CG method," International Journal of Infrared and Millimeter Waves, vol. 21, no. 9, pp. 1541-1560, 2000.
[23] F. H. Pereira and S. I. Nabeta, "Wavelet-based algebraic multigrid method using the lifting technique," Journal of Microwaves and Optoelectronics, vol. 9, no. 1, pp. 1-9, 2010.
[24] J. M. Ford, K. Chen, and N. J. Ford, "Parallel implementation of fast wavelet transforms," Numerical Analysis 39, Manchester University, Manchester, UK, 2001.
[25] H. J. Sips and H. X. Lin, High Performance Computing Course MPI Tutorial, Delft University of Technology Information Technology and Systems, Delft, The Netherlands, 2002.
[26] F. Alessandri, M. Chiodetti, A. Giugliarelli et al., "The electric-field integral-equation method for the analysis and design of a class of rectangular cavity filters loaded by dielectric and metallic cylindrical pucks," IEEE Transactions on Microwave Theory and Techniques, vol. 52, no. 8, pp. 1790-1797, 2004.
[27] T. A. Davis, "Algorithm 849: a concise sparse Cholesky factorization package," ACM Transactions on Mathematical Software, vol. 31, no. 4, pp. 587-591, 2005.
[28] D. A. Wood and M. D. Hill, "Cost-effective parallel computing," Computer, vol. 28, no. 2, pp. 69-72, 1995.

