# On the Influence of Partitioning Schemes on the Efficiency of Overlapping Domain Decomposition Methods 

P. Ciarlet, Jr.<br>CEA<br>CEL-V, D.MA, MCN<br>94195 Villeneuve-St-Georges Cedex France

F. Lamour

Cisi
Branche CEA et Défense
BP 28, 91192 Gif-sur-Yvette Cedex France

B. F. Smith

Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60439-4844
USA


#### Abstract

One level overlapping Schwarz domain decomposition preconditioners can be viewed as a generalization of block Jacobi preconditioning. The effect of the number of blocks and the amount of overlapping between blocks on the convergence rate is well understood. This paper considers the related issue of the effect of the scheme used to partition the matrix into blocks on the convergence rate of the preconditioned iterative method. Numerical results for Laplace and linear elasticity problems in two and three dimensions are presented. The tentative conclusion is that using overlap tends to decrease the differences between the rates of convergence for different partitioning schemes.


## 1 Introduction

One level overlapping Schwarz domain decomposition methods are simple, natural preconditioners for the parallel solution of the linear systems that arise from the discretization of elliptic PDEs. To apply these methods efficiently in parallel, one must partition the unknowns among the processors to achieve both load balancing and reduce required interprocessor communication. However, for block Jacobi methods the particular partition used can have a large effect on the numerical convergence rate.

Farhat and Lesoinne [11] and Farhat and Simon [12] addressed the problem of load balancing and communications between processors on a parallel architecture, and have run some experiments on parallel machines in order to compare the execution times. In this paper we study the effects of several partitioning strategies (with and without overlap) on the load balancing and communications but we focus on the con-
vergence rate of the method or, equivalently, on the number of iterations required to reach a prescribed tolerance.

We consider the numerical solution of elliptic PDEs in both two and three dimensions on unstructured grids using the finite element method. The resulting linear system is solved using a domain decomposition preconditioned Krylov space method designed for parallel computing. All of our numerical experiments are performed on sequential machines because we are mainly interested in the effect on the numerical convergence rate. Future studies will focus on the non-numerical effects introduced by parallel machines.

The paper is organized as follows. In Section 2, we present the elliptic problems to be solved, that is the equations and the boundary conditions. We also include a brief overview of the discretization, the iterative domain decomposition methods used in the paper. In the next Section, we give a detailed description of the partitioning heuristics, the influence of which we compare in Section 4.

## 2 Solving the elliptic problems

### 2.1 The problems and their discretization

Let $\Omega$ denote the domain imbedded in $\mathbb{R}^{2}$ or $\mathbb{R}^{3}$ and $\Gamma=\Gamma_{0} \bigcup \Gamma_{1}$ its boundary. The elliptic problems to be solved are either the Poisson problem with Dirichlet boundary conditions on $\Gamma_{0}$ and homogeneous Neumann boundary conditions on $\Gamma_{1}$ or the equations of linear elasticity in the domain $\Omega$ with Dirichlet boundary conditions on $\Gamma$. Let $u$ denote the (scalar) solution of the Poisson problem and $\underline{u}$ the (vector) solution of the elasticity equations. Finally, let $E$, Young's modulus, and $\nu$, Poisson's ratio, be two posi-
tive constants. Then the problems can be written as:

$$
\begin{gathered}
-\Delta u=f \text { in } \Omega, u=g \text { on } \Gamma_{0}, \frac{\partial u}{\partial n}=0 \text { on } \Gamma_{1}, \\
-\frac{E}{1+\nu}\left\{\Delta \underline{u}+\frac{1-\nu}{1-2 \nu} \nabla \nabla \cdot \underline{u}\right\}=\underline{f} \text { in } \Omega, \underline{u}=\underline{g} \text { on } \Gamma .
\end{gathered}
$$

In order to derive a discretization of these problems, the equations are replaced by a suitable variational formulation and then the finite element method [2] with the P1 finite element is used on the triangulations. Even though the value of the solution is known on $\Gamma_{0}$ via the Dirichlet boundary condition, we have chosen to keep these (trivial) equations in the linear system. Because of this, the unknowns are the approximations of the value of the true solution at all vertices of the triangulation, and the matrix is nonsymmetric.

### 2.2 Iterative and DD methods

The resulting nonsymmetric linear system is solved using the preconditioned bi-conjugate gradient stabilized method (Bi-CGSTAB), due to van der Vorst [23]. The preconditioner considered in this paper is designed with the help of domain decomposition principles.

In recent years, many papers have been devoted to the study of new domain decomposition methods. One of the main reasons for this renewed interest is the emergence of parallel and massively parallel computers. This is because many domain decomposition methods are naturally parallel. The additive Schwarz method we consider in this paper is very representative. This method was introduced by Dryja in [8] and Dryja and Widlund in [9]. For some results obtained on parallel architectures with this method, we refer the reader to [17].

Here, we consider the case of the one-level additive Schwarz preconditioner, that is local problems are solved without a coarse problem. It is well known that the addition of a coarse solver greatly reduces the condition number and therefore the number of iterations, but a coarse solver is defined from a coarse triangulation of the subdomain which may be costly to compute. In order to get a better condition number, it is possible, on the other hand, to consider overlapping subdomains. Here, the overlapping subdomains are simply obtained by extending the nonoverlapping partitions, at little additional cost. In the following, we thus allow subdomains to overlap.

Finally, let us briefly mention how the local problems are solved: the (complete) LU-factorization of
each local problem is computed and then used as a solver at every iteration of the method. To minimize the fill (number of non-zero entries) during the factorization, the matrix columns and rows are reordered with nested dissection following George and Liu's Sparsepack [15].

## 3 The partitioning heuristics

In order to partition the domain, the triangulation can be considered as a graph by identifying the notions of vertices and edges. We call this graph a finite element graph. Let $G=(V, E)$ be a graph where $V$ is the set of nodes and $E$ is the set of edges. Let $p$ be an integer, then we define the $p$ edge-partition of $G$ as a partition of $V$ into $p$ disjoint sets $V_{1}, V_{2}, \ldots V_{p}$ such that:

- $\bigcup_{i=1}^{p} V_{i}=V$,
- $\left|V_{i}\right| \approx\left|V_{j}\right|$ for $1 \leq i, j \leq p$ and $i \neq j$,
- $m=\mid\left\{\left(v_{i}, v_{j}\right) \in E\right.$ and $v_{i} \in V_{i}, v_{j} \in V_{j}$ with $i \neq$ $j\} \mid$ is as small as possible.

Besides the requirements of the $p$ edge-partitioning problem it may be necessary to impose that each resulting subset has to be connected.

For all the formulations and variants of graph partitioning, the general problem is a NP-hard problem. Many heuristics have been developed while few optimal results have been proved.

The heuristic methods proposed in the literature for finite element graphs can be roughly arranged in two categories: the greedy methods and the recursive methods. A detailed study of the heuristics presented here as well as a fair comparison of the corresponding partitions are given in [6].

Let us first introduce some definitions and notations about graphs that will be used in the following.

Let $N=|V|$ and $M=|E|$ be the number of nodes and edges respectively. The degree of a node $v$, denoted $d(v)$, is the number of edges that have one endpoint in $v$, i.e. the number of neighbors of $v$. Moreover we use $d$ to denote the average degree of a node in $G$.

The Laplacian matrix of a graph $G$, denoted $L(G)=\left(l_{i j}\right)_{i, j=1 \ldots N}$, is defined by:

$$
l_{i j}= \begin{cases}-1 & \text { if }\left(v_{i}, v_{j}\right) \in E \\ d\left(v_{i}\right) & \text { if } i=j \\ 0 & \text { otherwise }\end{cases}
$$

As we study finite element graphs we shall talk about the boundary of the graph that is clearly defined with respect to the meshes.

### 3.1 Recursive methods

### 3.1.1 Spectral bisection

Spectral partitioning methods are based on a particular eigenvalue, and on its associated eigenvector, of the Laplacian matrix of the graph to be partitioned.

In order to understand the intuitive justifications of these methods we must summarize some known properties of the eigenvalues of the Laplacian matrix of a graph and display their links with the graph bisection, or 2 edge-partitioning, problem.

Let $L(G)$ be the Laplacian matrix of $G$ and $N \geq 2$. Let $\lambda_{1} \leq \lambda_{2} \leq \lambda_{3} \leq \ldots \leq \lambda_{N}$ be the eigenvalues of $L(G)$.

To begin with, $\lambda_{1}$ is always equal to zero. The second smallest eigenvalue, $\lambda_{2}$, is also known as the algebraic connectivity of $G$. Let us recall that $\lambda_{2}=0$ if and only if $G$ is a disconnected graph. In other words, the number of eigenvalues equal to zero gives exactly the number of connected components of the graph. Let us assume, without loss of generality, that $\lambda_{2}$ is the first non zero eigenvalue.

The 2 edge-partitioning problem can be reformulated in the following terms. Let us consider a vector $q \in Q=\left\{\left(q_{i}\right) \in R^{N}, q_{i}= \pm 1, \sum q_{i}=0\right\}$ and the induced balanced partition for which a node $i$ is assigned to subset $V_{1}$ if $q_{i}=+1$ or to subset $V_{2}$ if $q_{i}=-1$.

The number of intersubset edges (or edge cuts) is equal to

$$
m=\frac{1}{4}(q, L(G) q)
$$

So the 2 edge-partitioning problem amounts to finding a vector $q \in Q$ which minimizes this quantity.

By relaxing this discrete minimization problem to a continuous problem, that is to the search of a vector $x$ such that $\|x\|_{2}^{2}=N$ and $\sum x_{i}=0$, one finds that the minimum of $(x, L(G) x)$ is obtained for $x=x_{2}$, where $x_{2}$ is the eigenvector associated with $\lambda_{2}$, named the Fiedler vector, [14].

Knowing this, the idea is to compute a vector $q \in Q$ by using $x_{2}$ in the following way. Let $x_{l}$ be the median value of the components of $x_{2}$. Then $q_{i}$ is defined as:

$$
q_{i}= \begin{cases}+1 & \text { if }\left(x_{2}\right)_{i}>x_{l} \\ -1 & \text { if }\left(x_{2}\right)_{i}<x_{l} \\ \pm 1 & \text { if }\left(x_{2}\right)_{i}=x_{l} \text { (for the balancing) }\end{cases}
$$

The partition induced by such a vector $q$ is known as the median cut partition. Of course this is not an
optimal partition. Although we cannot say how close to the optimal the median cut is, there are still some interesting properties about it.

First, there is an important result about the connectivity of a median cut partition which has been stated by Fiedler [14]. Briefly, if there are exactly $\frac{N}{2}$ strictly positive components and $\frac{N}{2}$ strictly negative components then both balanced subsets of the partition are connected. Regardless, the connectivity of one of the subsets is always guaranteed.

Secondly, it have been proved by Ciarlet, Chan and Szeto in [3] that for all $p \in Q,\left\|x_{2}-p\right\| \geq \| x_{2}-$ $q \|$. This result, which does not allow us to conclude about the optimality of $q$, can nevertheless reassures the promoters of the median cut method by assuming that it is a close enough choice.

### 3.1.2 Recursive spectral bisection

In order to partition a given graph into any number of subsets that is a power of two, the median cut partitioning method has been used as a step in a divide and conquer process, as in the Recursive Spectral Bisection algorithm ( $R S B$ ) due to Simon [22]. There the median cut algorithm is recursively applied to each subgraph induced by the bisection previously computed until the required number of subsets is obtained.

Moreover, as the computation of the Fiedler vector is time consuming, many studies have been devoted to speeding up the calculation of this particular vector. Barnard and Simon, [1], have accelerated the RSB algorithm by approximating the Fiedler vector. For this, they contract some edges in the graph in order to obtain a smaller graph and repeat this operation a certain number of times until the contracted graph is small enough. Then, they compute the Fiedler vector of the smallest graph. From the smallest graph, an approximation of the Fiedler vector of the previous graph is deduced by an interpolation technique. This approximated vector is then used as a starting point in an iterative method that computes the Fiedler vector. Once the refined Fiedler vector is obtained the process goes back to the larger graph and recomputes a Fiedler vector for this graph using the same strategy (interpolation and refinement), and so on. The algorithm is given below and it will be referred to as "RP $1 "$.

## Algorithm RP 1

1. Compute the Fiedler vector for the graph by:
(a) Constructing a series of smaller graphs $\left(G^{l}\right)_{l=1, \cdots, L}$ obtained by some contraction
operations applied to the original graph, (contraction step).
(b) Computing the Fielder vector for the smallest graph $G^{L}$.
(c) Constructing a series of Fiedler vectors corresponding to the series of graphs by:
i. interpolating the previously found Fiedler vector to the next larger graph in a way that provides a good approximation to next Fielder vector (interpolation step),
ii. computing from the given approximated Fielder vector, a more accurate vector (refinement step).
2. Sort vertices according to the size of entries in Fielder vector.
3. Assign half of the vertices to each subdomain.
4. Repeat recursively (divide and conquer).

The complexity of RP 1 is estimated to $O\left(M \log _{2} p\right)$.

### 3.1.3 Recursive multilevel algorithm

In a similar manner, Hendrickson and Leland in [20] used a multilevel technique to partition a graph. They first reduce the size of the graph, and derive a series of smaller graphs by contracting the edges of the original graph until the size of the last graph is small enough. They partition the smallest graph using the median cut technique. Then, they reflect the partition when uncontracting the series of graphs. Moreover, to improve the quality of the partition (in term of balancing and of number of intersubset edges) they perform a local optimization method which exchanges nodes between the subsets of the partition. Finally they repeat the previous sequence of steps on each subset until the total number of subsets is obtained. The algorithm below is called "RP 2".

## Algorithm RP 2

1. Construct a series of smaller graphs obtained by contraction operations applied to the original graph.
2. Partition the smallest graph using the median cut method.
3. Propagate the partition by:
(a) Uncontracting the smallest graph.
(b) Reflecting back the partition to the uncontracted graph.
(c) Refining locally the partition using a local optimization method.
(d) Repeating steps (a), (b), (c), until the original graph.
4. Repeat recursively (divide and conquer).

The complexity of RP 2 is estimated to $O\left(M \log _{2} p\right)$.

### 3.2 Greedy algorithms

### 3.2.1 Principles

According to the $p$ edge-partitioning problem a greedy algorithm can be described as an algorithm that computes each subset $V_{i}$ by simply accumulating nodes when traveling through the graph. The problematical questions are only: how to start and how to stop?

The way of accumulating nodes in each subset is obvious from the graph structure of the problem. A starting node $v_{s}$ is chosen and marked. The accretion process is done by selecting and marking the unmarked neighbors of $v_{s}$, then the unmarked neighbors of the neighbors of $v_{s}$ and so on as long as the expected total number of nodes is not reached. This can be viewed as successively building fronts.

The way of choosing a starting node $v_{s}$ will clearly affect the shape of the final partition. It will also influence the communication scheme, i.e. the number of existing edges between different subsets of the partition.

In the same way, the manner that one chooses the prescribed number of nodes among all the candidate nodes of the last front contributes to the quality of the final partition.

Thus a greedy heuristic for solving the $p$ edgepartitioning problem can be defined roughly by iterating the following 3 steps:

1. Choose a "good" starting node $v_{s}$,
2. Build fronts,
3. Stop according to some tie-break strategy in case of multiple choices and mark all the chosen nodes.

At present, there are no theoretical results on the "goodness" of one starting node. Neither are there results on how good a tie-break strategy is. For those two points only intuitive guesses help to design $p$ partitioning problem heuristics. However, an obvious justification of using greedy heuristics for solving the $p$ partitioning problem is that they are inexpensive. We have shown in [5], that for the general
case the overall complexity of such an algorithm is $O\left(N \max \left(p, d, \log _{2}\left(\frac{N}{p}\right)\right)\right)$. Note that the greedy algorithm was first proposed specifically for partitioning meshes (and for parallel processing) by Farhat [10].

### 3.2.2 Front oriented algorithm

The next algorithm, presented in detail in [5], implements the principles of a greedy method as well as some other original features. First, this algorithm builds connected subsets. On the other hand, it does not always provide well balanced subsets. The subsets are constructed in a concentric way around the boundary of the graph. Finally, for each subset, in case of multiple choices between the nodes of the last front, the tie-break strategy chooses those which are linked to as few unmarked nodes as possible.

More precisely, in the partitioning process, the starting node of each iteration $i$ is chosen in order to belong simultaneously to the boundary of $G$, to be an unmarked neighbor of a node of $V_{i-1}$ and to have a minimal positive current degree (in [10], the starting node is chosen among those which do not belong to the boundary of $G$ ). Here, the current degree of a node is the number of nodes connected to it which have not yet been selected, i.e. marked, during the accumulation step.

As for the tie-break strategy, it is achieved by keeping the nodes that have a minimal current degree. In fact, the algorithm does not simply build the subsets as mentioned but also check the connectivity of each of them. Whenever a subset is found to be multiconnected, the algorithm corrects the feature by reassigning small components to other subsets and by keeping the largest component. The algorithm "GP" is described next.

## Algorithm GP

1. If $i<p$
(a) Compute $n_{i}=\frac{N-\sum_{i=1}^{i-1} n_{j}}{p-(i-1)}$.
(b) Choose an unmarked node $v_{s}$ such that:
i. $v_{s}$ belongs to the current boundary,
ii. if the current boundary is not new, $v_{s}$ is a neighbor of a node belonging to $V_{i-1}$ (if possible ${ }^{1}$ ),
iii. $v_{s}$ has a minimal current degree.

Mark $v_{s}$ and initialize $V_{i}$ with $v_{s}$.

[^0](c) If there are unmarked neighbors of nodes of $V_{i}$, let $k$ be their number.
i. If $\left|V_{i}\right|+k<n_{i}$ then mark those nodes, add them to $V_{i}$ and update the current degree of their neighbors, and then return to 1.(c).
ii. Mark ( $\left.n_{i}-\left|V_{i}\right|\right)$ minimal current degree nodes and add them to $V_{i}$. Update the current and virtual boundaries.
Do $i=i+1$ and return to 1 .
(d) If there are no more unmarked neighbors of nodes of $V_{i}$ and if $\left|V_{i}\right|<n_{i}$ then unmark the nodes in $V_{i}$ and assign them to neighboring subsets. Return to 1.
2. Mark all the remaining nodes and add them into $V_{p}$. If $V_{p}$ is multiconnected then keep the largest component and unmark the nodes of the other components and assign these nodes to neighboring subsets.

The complexity of GP is $O(M)$.

### 3.3 Local optimization methods

Whatever partitioning methods may be used, many experiments have shown that local optimization methods can improve greatly the load balancing or the intersubset structure. One of the most well known optimization methods for improving a given partition is due to Kernighan and Lin [21].

Fundamentally the method is based on a given 2 edge-partition $\left(V_{1}, V_{2}\right)$ of the node set $V$ of a graph and tries to improve it by exchanging a subset of $V_{1}$ with one of $V_{2}$. The selection criterion of the subsets is determined from a gain function which is defined as follows. First, for $v \in V_{1}, g_{v}=d_{V_{2}}(v)-d_{V_{1}}(v)\left(d_{A}(v)\right.$ is the number of neighbors of $v$ which belong to $A$ ) and for $w \in V_{2}, g_{w}=d_{V_{1}}(w)-d_{V_{2}}(w)$. Then the gain obtained by exchanging a node $v \in V_{1}$ with a node $w \in V_{2}$ is equal to:

$$
g_{v, w}=g_{v}+g_{w}-2 \delta(v, w)
$$

where $\delta(v, w)$ is defined by:

$$
\delta(v, w)=\left\{\begin{array}{ll}
1 & \text { if }(v, w) \in E \\
0 & \text { otherwise }
\end{array} .\right.
$$

The authors of RP 1 ([1]), have chosen to use the KL algorithm as a postprocessing step after each bisection has been performed. Nevertheless, as the
complexity of the KL method is in the order of $O\left(N^{2} \log N\right)$, the KL method is actually invoked only when the number of nodes in each subset is less than 300 .

Because of the complexity of the KL algorithm, Fiduccia and Mattheyses [13] slightly modified the algorithm, and reduced the overall complexity to $O(M)$. Principally, the FM algorithm chooses to move one node after another instead of directly exchanging a pair of nodes. The authors of RP 2 generalized the FM method to an arbitrary number of subsets, [19] and [20]. In this case, the overall complexity is estimated to $O((p-1) M)$. This method is then invoked in step 3.(c) for refining the partition.

Finally, the authors of GP have designed, what they call a retrofitting method which has been suggested by the experiments conducted on their partitioning heuristic [4]. The first step tries to reshape the outlines of the subsets by deleting some excrescences. Generally these occur during the accumulation step of the partitioning process when some of the chosen nodes encounter prematurely another subset. Then some nodes are attached to the subset to which they belong by a single edge. By reassigning those nodes to neighboring subsets and by iterating the process until no more excrescences remain, the shape of the subsets is improved. Note that these boundary nodes are transferred to the subset which holds the highest number of its neighbors. It is not always possible to eliminate all the excrescences, so the reshaping step has to stop when the overall shape of the partition, i.e. the number of excrescences, does not seem to be improved over the iterations (more precisely five iterations). Here, one iteration consists of spanning the whole set of nodes $V$ and reassigning the excrescences.

Because of the reassignment of nodes to neighboring subsets in order to preserve the connectivity of the subsets, the resulting partition provided by GP is not balanced in most of the cases. Thus the second step of the retrofitting method looks for rebalancing the subsets by moving nodes from large subsets to small ones. One iteration of this consists of three parts. First, the largest and smallest subsets are determined. Then a set of nodes (or front) of the largest subset is reassigned to its smallest neighbor. Last, a set of nodes is reassigned to the smallest subset, coming from its largest neighbor. Moreover this process includes a connectivity test step which allows a move of a front only if the subsets of the new partition remain connected. The balancing process is iterated until the standard deviation approaches the optimal standard deviation
or until it does not decrease anymore.
The complexity of one iteration of the reshaping step is $O(M)$ while the complexity of the balancing step is in the order of $O(n d)$, where $n$ is the average number of node by subset. Note that when the retrofitting method is performed, neither the number of iterations for the reshaping steps nor for the balancing step are bounded a priori.

## 4 Numerical experiments

For our numerical experiments, we consider three domains. One is imbedded in $R^{2}$ and is the exterior of the section of a wing, made of four parts. The triangulation of this domain is due to T . Barth (from NASA Ames). The other two are imbedded in $\mathbb{R}^{3}$ : they are a domain bounded by two concentric spheres and an axle (part of an automobile). The first triangulation comes from L. Crouzet (from CEA) and the second one from M. Vidrascu (from INRIA). The three domains have been scaled in order to be in $[0,1]^{2}$ or $[0,1]^{3}$. The wing grid has 33677 nodes and 99520 edges, the spheres grid has 9020 nodes and 59418 edges and finally the axle grid has 25058 nodes and 156842 edges.

The Poisson problem is solved on the wing and the spheres, $f$ is set to $-9 x^{2} \sin 3 y+2 \sin 3 y$ in 2 D and 0 in $3 \mathrm{D} . \Gamma_{0}=\Gamma \bigcap\{(x, y), x<0.2\}$ and $g=x^{2} \sin 3 y$ in $2 \mathrm{D}, \Gamma_{0}=\Gamma$ and $g=x$ in 3 D . For the linear elasticity equations, solved on the axle, the constants are set to $E=1$ and $\nu=0.3 . \underline{f}$ is set to $\frac{30 E(1-\nu)}{(1+\nu)(1-2 \nu)}\left(\begin{array}{c}x^{4} \\ y^{4} \\ z^{4}\end{array}\right)$. Finally, $\underline{g}=\left(\begin{array}{l}x^{6} \\ y^{6} \\ z^{6}\end{array}\right)$.

The numerical experiments including both the partitioning heuristics and the iterative solver are tested through a common tool: the Portable Extensible Tools for Scientific computation (PETSc) developed by Gropp and Smith [16], [18]. PETSc is a software library for parallel and serial scientific computations. It provides a variety of packages and in particular the iterative solver. In addition to these existing packages, PETSc makes it easy to include any software written either in C or Fortran. The softwares corresponding to the partitioning heuristics have been given by their authors and have been used in PETSc as is. We have written interfaces for transferring inputs and parameters as well as for interpreting the results. The RP 1 code, written by Barnard and Simon, can be obtained
by a request to their authors ${ }^{2}$. The RP 2 code is copyrighted but can be obtained via a request to Hendrickson or Leland ${ }^{3}$. The GP code has been written by Ciarlet and Lamour.

The experiments were carried out on a Sun Viking SuperSparc.

### 4.1 The partitioning characteristics

The characteristics of the three partitioning heuristics can be summarized according to two parameters. The first one measures the balancing of the subsets. It gives a realistic idea of the difference of the size between large subsets and small ones according to the average size of the subsets: $\sigma / n$, where $n$ is the average number of nodes by subset and $\sigma$ is the standard deviation of the number of nodes by subset and is equal to $\sqrt{\frac{1}{p} \sum_{i=1}^{p}\left(n_{i}-n\right)^{2}}$. The second one measures the percentage of intersubset edges $(m / M)$.

From extensive numerical experiments [6], [7], we can see that $\sigma / n$ is very small in all cases (and optimal by definition for the recursive heuristics). But this is not enough to ensure the load balancing [7]. Further comments are made in the next Section.

The ratio $m / M$ measures the communication volume between the subdomains, where $M$ is given while $m$ is the parameter to be minimized by the heuristics. In general we can see that this volume is smallest for RP 2, followed closely by RP 1 . Nevertheless, GP gives reasonably close values [6], [7]. The impact of differences between communication volumes on iteration time is not measurable in our experiments as we have run them on a sequential machine.

### 4.2 The numerical characteristics

Basically, the numerical characteristics of the iterative method are the number of iterations needed to reach the prescribed tolerance of $10^{-5}$, and the amount of work by iteration. Let us first address the latter for nonoverlapping subdomains.

Mainly, an iteration consists in solving $p$ linear systems, one for each subdomain. Therefore the amount of work is proportional to the number of nonzero entries, or fill, of the LU-factorization of the local matrices, that is the restriction of the matrix to each subdomain. Concerning the average fills, GP gives generally better results than the recursive heuristics (for numerical results, we refer the reader to [7]). This can be

[^1]explained in part by the average fill of the local matrices before factorization. It is equal to $\frac{(M+N)-m}{p}$ and therefore smaller for GP. The reordering and factorization process does not seem to reduce these original differences. The reordering is done using Sparsepack's nested dissection routines.

Moreover, as these methods are designed for parallel architectures, it is natural to study the load balancing, that is not the average fill but the fill for each subdomain. The most important values are the minimum and the maximum fills for a given partition. In most cases, there is a great imbalance [7].

The maximum value is representative of the largest amount of work on a processor, in the case of a parallel experiment. Generally, these are very close for the three heuristics. As a function of the number of partitions $p$, RP 1 gives the best results for small $p$ whereas GP performs best for large $p$, both for overlapping and nonoverlapping subdomains [7].

Let us now investigate the possible origins of the imbalance. It can be caused by some original imbalance, i.e. the discrepancy in fill of the local matrices before they are factored. It can also be caused by the structure of these matrices which affects the reordering and factorization process. The second cause is not clearly related to the partitioning heuristic, but the first one is.

Indeed, the constraint, when partitioning, is on the number of vertices. Unfortunately, the nonzero entries of the local matrices correspond exactly to the edges of the graph. Thus, in order to have balanced local matrices, one has to replace the constraint on vertices by a constraint on edges. But it can not be easily done, since the average number of edges by subset $\frac{(M+N)-m}{p}$ is not known a priori: a value for $m$ is obtained only after the partition is computed.

Nevertheless, for this particular problem, it is possible to design special postprocessing strategies. Here, once $m$ is determined, some "retrofitting" technique could be applied very simply, targeting the number of edges instead of the number of vertices. Note that, in this case, the original constraint on the number of vertices is not crucial any more. Briefly, let us mention another way of avoiding this problem in 2 D . As a matter of fact, for a given region $R$ of the 2D-triangulation, if $T_{R}$ denotes the number of triangles and $M_{R}$ the number of edges in the region, then

$$
3 T_{R} \approx 2 M_{R}
$$

Based on this observation, it can be interesting to partition the triangulation in terms of triangles, as well balanced subsets of triangles correspond roughly to
well balanced subsets of edges. In 3D, however, there is no such relationship between the number of tetrahedra and edges in a region.

Finally, let us mention that Vanderstraeten, Zone, Keunings and Wolsey have proposed in [24] a postprocessing strategy to balance a quantity that approximates the actual fill.

Now, we consider the number of iterations as a function of both the number of subdomains $p$ and the overlap.

| $p$ | 16 | 32 | 64 | 128 | 256 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| RP 1 | 146 | 158 | 134 | 136 | 152 |
| RP 2 | 83 | 94 | 107 | 142 | 79 |
| GP | 105 | 83 | 96 | 119 | 144 |

Wing: Number of iterations without overlap.

| $p$ | 16 | 32 | 64 | 128 | 256 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| RP 1 | 21 | 25 | 22 | 28 | 27 |
| RP 2 | 19 | 21 | 21 | 29 | 22 |
| GP | 19 | 21 | 27 | 24 | 29 |

Spheres: Number of iterations without overlap.

| $p$ | 32 | 64 | 128 | 256 |
| :---: | :---: | :---: | :---: | :---: |
| RP 1 | 11 | 13 | 14 | 17 |
| RP 2 | 10 | 13 | 16 | 16 |
| GP | 15 | 16 | 17 | 20 |

Axle: Number of iterations without overlap.
First, we consider the nonoverlapping case. It is the only case we calculate for the axle because of memory limitations. We can see from the Tables that no heuristic performs really better than the others in terms of the number of iterations. Moreover, these numbers are relatively close to one another for a given $p$, especially when they are small (see the 3D examples). Finally, for the spheres and the axle we note that the number of iterations increases only slowly with the number of subdomains. This is particularly true for the axle.

Figures 1 to 6 illustrate the behavior of the number of iterations according to the overlap, the number of subdomains being fixed. For Figures 5 and 6, the missing parts of the curves are due to memory limitations. In the case of overlapping subdomains we can also see that no heuristic performs dramatically better. The different values of the number of iterations tend to be more clustered as the overlap increases. Once again, the behavior of the curves is not strictly monotonic. Yet, it tends to decrease as the overlap increases.

We note that the reason for the larger number of iterations for the problem in two dimensions is due to the Neumann boundary conditions which usually decrease the convergence rate of block Jacobi type preconditioners.

## 5 Conclusion

We solved some elliptic PDEs in two and three dimensions on unstructured grids using the one level overlapping Schwarz method. We partitioned the triangulations with three different heuristics and studied the impact of the partitioning on the numerical results. Based on these experiments, it appears that the three heuristics behave very similarly in terms of load balancing. For the convergence rate the increasing amount of overlap tends to smooth out the differences between the iterative solvers. For nonoverlapping subdomains it is difficult to point out any advantage of one heuristic over the others, since different heuristics perform best on different problems.

## Acknowledgments

We thank Horst Simon for providing the code of RP 1 and Bruce Hendrickson and Robert Leland for allowing us to use the RP 2 code.

The work of the first author was partially supported by the DGA/DRET under contract 93-1192. The work of the second author was partially supported by the National Science Foundation under contract ASC 92-01266, the Army Research Office under contract DAAL03-91-G-0150, and ONR under contract ONR-N00014-92-J-1890. The work of the third author was supported by the Office of Scientific Computing, U.S. Department of Energy, under contract W-31-109-Eng38.

## References

[1] S.T. Barnard and H.D. Simon. A fast multilevel implementation of recursive spectral bisection for partitioning unstructured problems. Technical report, NASA Ames Research Center, RNR-092033, 1992.
[2] P. Ciarlet. The finite element method for elliptic problems. North Holland, 1978.
[3] P. Ciarlet, Jr, T.F. Chan, and W.K. Szeto. On the optimality of the median cut spectral bisec-
tion graph partitioning method. Technical report, UCLA, CAM 93-14, 1993.
[4] P. Ciarlet, Jr and F. Lamour. A front-oriented approach to partitioning sparse graphs. Technical report, in preparation.
[5] P. Ciarlet, Jr and F. Lamour. An efficient low cost greedy graph partitioning heuristic. Technical report, UCLA, CAM 94-1, 1994.
[6] P. Ciarlet, Jr and F. Lamour. Recursive partitioning methods and greedy partitioning methods: a comparison on finite element graphs. Technical report, UCLA, CAM 94-9, 1994.
[7] P. Ciarlet, Jr, F. Lamour, and B.F. Smith. On the influence of the partitioning schemes on the efficiency of overlapping domain decomposition methods. Technical report, UCLA, CAM 94-23, 1994.
[8] M. Dryja. An additive Schwarz algorithm for two- and three-dimensional finite element elliptic problems. In T. Chan, R. Glowinsky, J. Périaux, and O. Widlund, editors, Domain Decomposition Methods. SIAM, Philadelphia, PA, 1989.
[9] M. Dryja and O. B. Widlund. An additive variant of the Schwarz alternating method for the case of many subregions. Technical report, 339, Dept of Computer Science, Courant Institute, 1987.
[10] C. Farhat. A simple and efficient automatic FEM domain decomposer. Computers and structures, vol $28, \mathrm{n}^{\circ} 5$, pp 579-602, 1988.
[11] C. Farhat and M. Lesoinne. Automatic partitioning of unstructured meshes for the parallel solution of problems in computational mechanics. International Journal for Numerical Methods in Engineering, vol 36, n ${ }^{\circ}$ 5, pp 745-764, 1993.
[12] C. Farhat and H.D. Simon. TOP/DOMDEC- a software tool for mesh partitioning and parallel processing. Technical report, NASA Ames Research Center, RNR-93-011, 1993.
[13] C.M. Fiduccia and R.M. Mattheyses. A lineartime heuristic for improving network partitions. Proceedings of the 19th IEEE Design Automation Conference, IEEE, pp 175-181, 1982.
[14] M. Fiedler. Algebraic connectivity of graphs. Czechoslovak Mathematical Journal, 23, (98), pp 298-305, 1973.
[15] A. George and J. Liu. Computer solution of large sparse positive definite systems. Prentice-Hall, Englewood Cliffs, NJ, 1981.
[16] W.D. Gropp and B.F. Smith. Portable, Extensible Toolkit for Scientific Computation (PETSc). Available at info.mcs.anl.gov in the directory pub/pdetools via anonymous ftp.
[17] W.D. Gropp and B.F. Smith. Experiences with domain decomposition in three dimensions: overlapping Schwarz methods. In A. Quarteroni, Y.A. Kuznetsov, J. Périaux, and O. Widlund, editors, Domain Decomposition Methods in Science and Engineering. AMS contemporary mathematics, vol 157, 1993.
[18] W.D. Gropp and B.F. Smith. Scalable, extensible, and portable numerical libraries. Proceedings of the Scalable Parallel Libraries Conference, IEEE, pp 87-93, 1993.
[19] B. Hendrickson and R. Leland. Multidimensional spectral load balancing. Technical report, SAND 93-0074, 1993.
[20] B. Hendrickson and R. Leland. A multilevel algorithm for partitioning graphs. Technical report, SAND 93-1301, 1993.
[21] B.W. Kernighan and S. Lin. An efficient heuristic for partitioning graphs. The Bell System Technical Journal, vol 49, no 2, pp 291-307, 1970.
[22] H.D. Simon. Partitioning of unstructured problems for parallel processing. Computing Systems in Engineering, vol 2, n ${ }^{\circ} 2 / 3$, pp 135-148, 1991.
[23] H. A. van der Vorst. Bi-cgstab: a fast and smoothly converging variant of bi-cg for the solution of nonsymmetric linear systems. SIAM J. Sci. Stat. Comp., vol 13, n ${ }^{\circ}$ 2, pp 631-644, 1992.
[24] D. Vanderstraeten, O. Zone, R. Keunings, and L. A. Wolsey. Non-deterministic heuristics for automatic domain decomposition in direct parallel finite element calculations. In R.F. Sincovec, editor, in Parallel Processing for Scientific Computing. SIAM, 1993.


[^0]:    ${ }^{1}$ It may not be possible to find a node neighboring the previously built subset if the boundary is multiconnected.

[^1]:    ${ }^{2}$ simon@nas.nasa.gov or barnard@nas.nasa.gov.
    ${ }^{3}$ bahendr@cs.sandia.gov or rwlelan@cs.sandia.gov.

