On the validity of a front-oriented approach to partitioning large sparse graphs with a connectivity constraint

P. Ciarlet, Jr.*

Commissariat à l'Energie Atomique, LV, D.MA, MCN, F-94195 Villeneuve-St-Georges Cedex, France
and Department of Mathematics, University of California at Los Angeles, CA 90095, USA
E-mail: ciarlet@limeil.cea.fr, ciarlet@math.ucla.edu

F. Lamour**

Department of Mathematics, University of California at Los Angeles, CA 90095, USA
E-mail: lamour@math.ucla.edu, lamour@limeil.cea.fr

In this paper we consider the problem of partitioning large sparse graphs, such as finite element meshes. The heuristic which is proposed allows to partition into connected and quasi-balanced subgraphs in a reasonable amount of time, while attempting to minimize the number of edge cuts. Here the goal is to build partitions for graphs containing large numbers of nodes and edges, in practice at least 10^4. Basically, the algorithm relies on the iterative construction of connected subgraphs. This construction is achieved by successively exploring clusters of nodes called fronts. Indeed, a judicious use of fronts ensures the connectivity of the subsets at low cost: it is shown that locally, i.e. for a given subgraph, the complexity of such operations grows at most linearly with the number of edges. Moreover, a few examples are given to illustrate the quality and speed of the heuristic.

1. Introduction

Graph partitioning methods are commonly used for a wide variety of problems: problem segmentation, circuit design, parallel computations, domain decomposition

methods and many others [11,14,19]. Basically, the global problem is decomposed into smaller subproblems which interact as little as possible with one another.

There are different formulations and many variants of the graph partitioning problem, depending on whether one looks at it from the vertex point of view or from the edge point of view. According to the edge point of view, on which we focus in this paper, the problem can be expressed mathematically in the following way, for a 2-edge partition or bisection. Let $G = (V, E)$ be a graph where V is the set of nodes and E is the set of edges. The 2-edge partition of G is a partition of V into two disjoint sets V_1 and V_2 such that:

- $V = V_1 \bigcup V_2$,
- $\max(|V_1|, |V_2|) \le \alpha |V|$,
- $|E_0|$ is as small as possible, where $E_0 = E \cap V_1 \times V_2$ is the set of edge cuts or intersubset edges.

Here, we use the classical notation $|.|$ for the cardinality of a set. The value of the constant α ($1/2 \le \alpha < 1$) characterizes the looseness in the balancing of the partition. It is well known that the previous problem is NP-complete for any value of the parameter α. Thus, many heuristics have been designed to approximate as closely as possible this bisection problem. Among them, node interchanging methods [15,24], branch and bound methods [7,29], maximum flows [6], maximum matching [25] or eigenvalue computations [20,27].

In the general case, the p-edge partition of G, for $p \ge 2$, is a partition of V into p disjoint subsets $V_1, V_2, \ldots, V_p$ such that:

- $V = \bigcup_{i=1}^{p} V_i$,
- $|V_i| \approx |V_j|$,
- $|E_0|$ is as small as possible, where $E_0 = E \cap \bigcup_{i \ne j} V_i \times V_j$.

This more general problem is also NP-complete, and many heuristics have likewise been constructed to approximate it. First, there is an important class of such methods which rely on one of the bisection heuristics of the previous paragraph. Initially, the graph is partitioned into two parts and then each subgraph is divided into two, and the process goes on recursively. Among them, eigenvalue computations [2,21] or multilevel techniques [23]. Many other heuristics have been designed to deal directly with the general p-edge partition problem, such as linear programming transportation [3], greedy algorithms [13], node interchanging methods [4] or geometrical approaches [1,26].

Note that some of these methods can be utilized jointly. Particularly, node interchanging methods require some initial partition. On the other hand, these methods are often used as a postprocessing step, i.e. after a first approximation to the solution of the p-edge partition problem has been determined, because they generally increase the quality of the partition.

In the following, we give a theoretical framework to the greedy method introduced in three earlier technical reports [8–10]. In addition, a new load-balancing postprocessing step is presented. Both methods are based on

particular clusters of nodes which we call *fronts*. The next section recalls some graph definitions and notations together with the *p*-edge partitioning and load-balancing problems. Section 3 is devoted to fronts and connectivity in graphs. In sections 4 and 5, the partitioning heuristic and the postprocessing steps are described thoroughly. In the next section, some experiments are listed on a number of graph examples. Finally, we summarize the main results in section 7.

2. Notations and problems

First, let us recall a few definitions about graphs and set the notations that will be used throughout this paper. For more details, we refer the reader to [5]. Let $G = (V, E)$ be an *undirected* graph where V is the set of nodes and E the set of edges. Only *simple* graphs will be considered hereafter, i.e. without multiple edges or loops. Let $N = |V|$ and $M = |E|$ be the number of nodes and edges respectively.

G is said to be *connected* if, for any pair of nodes (v, w), either (v, w) is an edge or there exist k nodes (x_i) such that the edges of the *path* $(v, x_1), (x_1, x_2), \ldots, (x_{k-1}, x_k)$, (x_k, w) belong to E. Otherwise, G is said to be *multiconnected*.

For a node v, let $\Gamma(v)$ denote the set of its *neighbors*, that is the set of nodes w such that (v, w) belongs to E, and let $\underline{\Gamma}(v) = \Gamma(v) \cup \{v\}$. More generally, for a nonempty subset W of V, $\underline{\Gamma}(W) = \cup_{w \in W} \underline{\Gamma}(w)$ and, by extension, $\underline{\Gamma}^l(W) = \underline{\Gamma}(\underline{\Gamma}^{l-1}(W))$ for $l \geq 2$ ($\underline{\Gamma}^1(W) = \underline{\Gamma}(W)$). For practical reasons, we define $\underline{\Gamma}^0(W) = W$. The *degree* of a node v, denoted by $d(v)$, is $|\Gamma(v)|$. For a node v and a subset W, let $d_W(v)$ be the number of neighbors of v belonging to W. Finally, let $\bar{d}$ be the average degree of a node in G. By definition, $\bar{d} = 2M/N$.

For any pair of nodes (v, w), the *distance* $\delta_G(v, w)$ corresponds to the length (i.e. number of edges) of the shortest path linking v to w in G ($\delta_G(v, v) = 0$); by extension, if there is no path linking two nodes v and w in the graph, then $\delta_G(v, w) = +\infty$ [5]. For a nonempty subset W and a node v, let $\delta_G(W, v) = \min_{w \in W} \delta_G(w, v)$. Finally, $\max_{v,w} \delta_G(v, w)$ is called the *diameter* of the graph.

In the following, as we partition graphs that come from physical meshes, usually two- or three-dimensional, we talk about the *boundary* of the graph. Let n_b be the number of boundary nodes.

For the partitioning problem considered here, we shall denote by p the number of subsets V_i and $n_i = |V_i|$. Let $\bar{n}$ be the average of (n_i), that is, $\bar{n} = \lfloor N/p \rfloor$[1]. Moreover, for $1 \leq i \leq p$, let $E_i = E \cap V_i \times V_i$ and $G_i = (V_i, E_i)$ and let $E_0 = E \cap \bigcup_{i \neq j} V_i \times V_j$. Note that $E_0, E_1, \cdots E_p$ is a disjoint partition of E. If we let $m_j = |E_j|$, then the *p*-edge partitioning problem we solve hereafter is:

Find *p* *disjoint* subsets $V_1, V_2, \cdots, V_p$ of V such that

- $V = \bigcup_{i=1}^{p} V_i$,

---

[1] $\lfloor . \rfloor$ stands for the lower integer part of a number.

- $n_i \approx n_j$,
- the subgraphs G_i are *connected*,
- m_0 is minimized.

Because of the added connectivity constraint, it often happens that the subsets that are obtained, although connected, are not balanced. Thus a postprocessing heuristic, based on reassigning nodes or sets of nodes is used afterwards. Two load balancing steps, called LB1 and LB2, aim at decreasing the value of the standard deviation of (n_i), i.e. $\sigma = \sqrt{(1/p)\sum_{i=1}^{p}(n_i - \bar{n})^2}$. LB1 relies on successive *single* node reassignments, enforcing the connectivity constraint on the subgraphs and trying to minimize the number of edge cuts m_0. LB2, on the other hand, allows for successive *multiple* nodes reassignments (by fronts), enforcing only the connectivity constraint.

3. Fronts in graphs

There are two obvious ways of exploring a graph G, from a given vertex v:

- examine all its neighbors and then pass to one of them, or
- examine one of its neighbors w and then pass to w.

The first approach is called *breadth-first search* and the second *depth-first search* (see, for instance, [17,18]). In this paper, we consider only algorithms based on breadth-first-like search. The idea is to explore the graph by *fronts*. From a non-empty subset W of V (particularly singletons), examine all nodes at distance 1, then all nodes at distance 2, etc. Fronts at distance k, $k \geq 0$, of W are defined by

$$F(W,k) = \{v \in V, \delta_G(W,v) = k\}.$$

Proposition 1
For a nonempty subset W of V and $k \geq 1$, $\underline{\Gamma}^k(W) = \underline{\Gamma}^{k-1}(W) \cup F(W,k)$ and $\underline{\Gamma}^{k-1}(W) \cap F(W,k) = \emptyset$.

Proof
The proof proceeds by induction.
    For $k = 1$, by definition,

$$
\begin{aligned}
\underline{\Gamma}(W) = \bigcup_{w \in W}\underline{\Gamma}(w) &= \bigcup_{w \in W}\Gamma(w) \cup W \\
&= \{v \in V, \exists w \in W, (v,w) \in E\} \cup W \\
&= \{v \in V \backslash W, \exists w \in W, (v,w) \in E\} \cup W \\
&= \{v \in V, \delta_G(W,v) = 1\} \cup W \\
&= F(W,1) \cup \underline{\Gamma}^0(W). \quad\quad\quad\quad (1)
\end{aligned}
$$

In addition, $\underline{\Gamma}^0(W) = W$ and $F(W, 1) = \{v \in V \backslash W, \ \delta_G(W, v) = 1\}$: they are disjoint.

Now, let us assume that the result is true for $l \leq k - 1$ and prove it also holds for k. First, by assumption, we have

$$\underline{\Gamma}^{k-1}(W) = \underline{\Gamma}^{k-2}(W) \cup F(W, k-1) = \cdots = \cup_{l=0}^{k-1} F(W, l)$$
$$= \{v \in V, \delta_G(W, v) \leq k - 1\}. \tag{2}$$

Thus $\underline{\Gamma}^{k-1}(W) \cap F(W, k) = \emptyset$. If we use (1) and replace W by $\underline{\Gamma}^{k-1}(W)$, then we have

$$\underline{\Gamma}(\underline{\Gamma}^{k-1}(W)) = F(\underline{\Gamma}^{k-1}(W), 1) \cup \underline{\Gamma}^0(\underline{\Gamma}^{k-1}(W)), \ \text{or}$$

$$\underline{\Gamma}^k(W) = F(\underline{\Gamma}^{k-1}(W), 1) \cup \underline{\Gamma}^{k-1}(W).$$

To complete the proof, one has to check that $F(\underline{\Gamma}^{k-1}(W), 1) = F(W, k)$:

$$F(\underline{\Gamma}^{k-1}(W), 1) = \{v \in V \backslash \underline{\Gamma}^{k-1}(W), \exists z \in \underline{\Gamma}^{k-1}(W), \delta_G(z, v) = 1\}$$
$$= \{v \in V, \delta_G(W, v) > k - 1, \exists z, \delta_G(W, z) \leq k - 1, \delta_G(z, v) = 1\} \quad \text{by (2)}$$
$$= \{v \in V, \delta_G(W, v) = k\}.$$

$\square$

By definition, once a node has been reached and processed, it is *marked*. Then, by rewriting Proposition 1 in the following way, $F(W, k) = \underline{\Gamma}^k(W) \backslash \underline{\Gamma}^{k-1}(W)$, one gets a natural way of traveling simply through the graph by successively exploring fronts, where each new front is the set of neighbors which have not been marked yet: this process is called an *expansion* by fronts in G from W. As for breadth-first search (and depth-first search), the complexity of exploring the graph by fronts is $O(M)$, as each edge is considered exactly twice during the process.

We shall use this method by fronts in the next section to build a partitioning heuristic. Now, as our goal is to obtain connected subsets of the set of nodes, let us prove a few simple results regarding connectivity: more precisely how to check that a graph or a subgraph is connected with the help of fronts.

Proposition 2

A graph $G = (V, E)$ is connected if and only if either

1. there exists a connected subset W_c of V such that the expansion by fronts from W_c spans V, or
2. for all subsets W of V, the expansion by fronts from W spans V.

Proof
  (i)   If G is connected, then all distances are finite and 2 is true.
 (ii)   If 2 is satisfied, then 1 is clearly verified.
(iii)   If 1 holds, then let us check that G is connected. Let (v_1, v_2) be a pair of nodes.

As the expansion by fronts from W_c spans V, there exist two nodes (w_1, w_2) belonging to W_c and two paths in G linking v_1 to w_1 and v_2 to w_2. By assumption, as W_c is connected, there exists a path linking w_1 to w_2. Thus there is a path in G linking v_1 to v_2: G is connected. $\square$

Now, let $G = (V, E)$ be a connected graph and W a subset of V. We give in the next proposition several ways of checking whether the subgraph with W removed is still connected. Let $\Gamma(W) = F(W, 1)$ and

- $V_- = V \backslash W, E_- = E \cap (V_- \times V_-), G_- = (V_-, E_-),$
- $G_\Gamma = (\Gamma(W), E_- \cap \Gamma(W) \times \Gamma(W))$ and
- s a vertex of $\Gamma(W)$.

In the next proposition, we give a few conditions to check whether G_- is connected or not.

Proposition 3
G_- is connected if either

1.   the expansion by fronts from s in G_- spans V_-, or
2.   the expansion by fronts from s in G_- spans $\Gamma(W)$, or
3.   the expansion by fronts from s in G_Γ spans $\Gamma(W)$.

Moreover, the connectivity of G_- implies conditions 1 and 2 but does not imply condition 3.

Proof
(i)   Clearly, the fact that G_- is connected implies condition 1. Also, condition 1 implies condition 2.
(ii) Let us prove that condition 2 yields the connectivity of G_-. Let (v, w) belong to $V_- \times V_-$. G is connected by assumption and therefore there exists a path $(x_1, \cdots, x_k)$ of nodes of V such that $(v, x_1), \ldots, (x_k, w)$ are in E. If no x_i belongs to W, then the path is in G_-. Now, if there is at least one x_i in W, let i_0 (resp. i_1) denote the first (resp. last) index such that $x_i \in W$. If we let $x_0 = v$ and $x_{k+1} = w$, then x_{i_0-1} and x_{i_1+1} belong to $\Gamma(W)$ and by assumption (condition 2), there exists a path $y_1, \cdots, y_l$ in G_- linking x_{i_0-1} to x_{i_1+1}. Then the path $x_1, \cdots, x_{i_0-1}, y_l, \cdots, y_l, x_{i_1+1}, \cdots, x_k$ is in G_- and links v to w: G_- is connected.
(iii) G_Γ is embedded in G_-, thus condition 3 implies condition 2.
(iv) Finally, to see that condition 2 does not yield condition 3, one can consider the example of the removal of a single node from an annulus (i.e. a node is exactly one endpoint of two edges) of size at least four. $\square$

The complexities $(C_i)_{i=1,2,3}$ of checking the connectivity of G_- by properties 1 to

3 of proposition 3 are $C_1 = C_2 = O(M_-)$ and $C_3 = O(M_\Gamma)$, and in any case $C_3 \leq C_2 \leq C_1$.

4. Greedy partitioning heuristics

Briefly there are two schools for computing a solution of the p-edge partitioning problem for large graphs: from the matrix point of view, i.e. using either the adjacency matrix or the Laplacian matrix [16] associated to the graph, or directly from the graph itself. The first point of view is based on spectral properties of one of these two matrices. Among the different works, let us mention those of Barnes [3], Rendl and Wolkowicz [28], Simon [30] and Hendrickson and Leland [23]. The other way is to work directly on the graph by applying greedy methods: the heuristic presented below is based on such a greedy algorithm. Let us mention that the first greedy algorithm for partitioning finite element meshes was proposed by Farhat [13].

Greedy algorithms are a natural and naive way to look at some problems. In some cases, they seem to give reasonable results, especially here for approximating the solution of the p-edge partitioning problem. In this case, a greedy algorithm can be described as an algorithm that computes one after the other each subset V_i by simply accumulating nodes when traveling through the graph. Evidently, we choose to travel in the graph by fronts. Then the only problematical questions are: how to start and how to stop?

A starting node v_s is chosen and marked. The accretion process is done by successively building fronts from v_s and so on as long as the expected total number of nodes is not reached.

The way of selecting a starting node v_s will clearly affect the shape of the final partition. It will also influence the communications, i.e. the number of existing edges between different subsets of the partition.

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

Thus, for each subset, one can summarize the greedy procedure by the following three steps:

1. select a "good" starting node v_s,
2. accumulate enough descendants of v_s,
3. stop according to some tie-break strategy in case of multiple choices.

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-edge partitioning heuristics. However, if we suppose that the tie-break strategy only involves *sorting*, then an obvious justification of using greedy heuristics is that they are inexpensive, as we can see in the following proposition.

Proposition 4
The overall complexity of a greedy algorithm applied to the p-edge partitioning problem is $O(M) + \{\max(p, \log_2(N/p))\}O(N)$.

Proof
By using the above procedure, one can easily see that the complexities of steps 1, 2 and 3 are respectively $O(pN)$, $O(M)$ and $O(p\bar{n}\log_2(\bar{n}))$, where $\bar{n}$ is the number of nodes per subset. $\square$

Remark 1
The optimal choice for the tie-break strategy is: select the required number of nodes among the candidates of the last front in such a way that the number of edge cuts is minimized, i.e. the number of edges between the current subset and the rest of the graph. It is clear that this optimal choice is NP-complete. Thus there is a need for an approximate choice based on sorting to reduce the complexity.

4.1. The GP heuristic

First, let us recall that, as we consider meshes, we can use the notion of the boundary of the graph.

As greedy algorithms do, our algorithm builds iteratively the different subsets of the partition by accumulating nodes in each subset and marking them once they have been selected. For the first iteration, the starting node is chosen among the boundary nodes with minimum degree. On the other hand, at each iteration one can define the *current boundary* as the set of nodes that belong to the boundary of the graph and that have not been marked yet. We can extend the notion of current boundary even when there are no more unmarked boundary nodes, by defining a new one and updating n_b, the number of boundary nodes. In that case the new current boundary is the set of unmarked nodes that are the neighbors of marked nodes. From there each new starting node is chosen among nodes which stand on the current boundary and are neighbors of the previously built subset. Moreover, this node is preferably chosen among the candidates with minimum *current degree*, i.e. the number of unmarked neighbors. Note that the current degree starts from the degree of the node and decreases to zero as the partitioning algorithm proceeds.

At some point, the current boundary can become empty. A way to redefine the current boundary at low cost is to introduce the notion of *virtual boundary*. Actually, this virtual boundary is defined, after each iteration, as the set of unmarked nodes that are neighbors of marked nodes, excluding current boundary nodes. Then, the new current boundary is defined as the virtual boundary, and the new virtual boundary is built anew.

The greedy partitioning algorithm "GP" is described next. For the moment, we consider the case of a connected graph G and prove results concerning its properties and its complexity. In subsection 4.4, we describe how to use GP on multiconnected graphs.

Let $i = 1$.

Step 1. If $i < p$

  (a) Compute

$$n_i = \left\lfloor \frac{N - \sum_{j=1}^{i-1} n_j}{p - (i-1)} \right\rfloor .$$

  (b) Select 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[2]),
    iii. v_s has a minimal current degree.
    Mark v_s and initialise V_i with v_s.
  (c) $K = 0$.
    If there are unmarked neighbors of nodes of V_i then
    i. let K be their number.
    ii. If $|V_i| + K < n_i$ then
        mark those nodes and add them to V_i,
        update the current degree of their neighbors,
        return to 1(c).
  (d) If $|V_i| + K \geq n_i$ then
        mark $(n_i - |V_i|)$ minimal current degree nodes and add them to V_i,
        update the current degree of their neighbors,
        update the current and virtual boundaries,
        do $i = i + 1$ and return to 1.
  (e) If there are no more unmarked neighbors of nodes of V_i then
        unmark the nodes in V_i and assign them to neighboring subsets,
        return to 1.

Step 2. Mark all the remaining nodes and add them into V_p. If V_p is multi-connected then keep the largest component and unmark the nodes of the other components and assign these nodes to neighboring subsets.

Remark 2

This heuristic can be used for any graph in a simpler version, that is, without the notion of boundary, be it current or virtual. Nevertheless, numerous numerical experiments show that, when compared to the original heuristic for physical meshes, this simpler version of the heuristic is slightly slower, and gives worse results in terms of edge cuts.

---

[2] It may not be possible to find a node neighboring the previously built subset if the boundary is multiconnected.

4.2. *Tie-break strategy*

Let us go into the details of our tie-break strategy when building a given subset V_{i_0}. Note that because of its intrinsic complexity, the optimal choice is impossible to determine. Suppose that the kth front $F(v_s, k)$ contains too many nodes $(v_i)_{i=1,\cdots,p_k}$ and that we want only q of them. The optimal choice corresponds to selecting q nodes such that the number of edges between them and nodes that do not belong to the current subset (i.e. $d(v_i) - d_{V_{i_0}}(v_i)$ for one node) is minimized. Then the objective is to find q nodes $(v_{\sigma(j)})$ minimizing

$$f_{opt}(v_{\sigma(j)}) = \sum_{j=1}^{q}\{d(v_{\sigma(j)}) - d_{V_{i_0}}(v_{\sigma(j)})\}.$$

If we let $d_c^l(v)$ be the current degree of a node v after the lth front is obtained, then $d_c^k(v) = d(v) - \sum_{i \leq i_0} d_{V_i}(v)$. Thus the objective can be reformulated as

$$f_{opt}(v_{\sigma(j)}) = \sum_{j=1}^{q}\{d_c^k(v_{\sigma(j)}) + \sum_{i < i_0} d_{V_i}(v_{\sigma(j)})\}. \tag{3}$$

Here, we make two modifications to get a simpler objective function. First, we can neglect the contribution coming from other subsets, as this contribution appears anyway in the total number of edge cuts. Indeed, a node v which has not been assigned yet necessarily belongs to V_i, for $i \geq i_0$. Then the edges corresponding to $\{\sum_{i < i_0} d_{V_i}(v)\}$ will be cut. In other words, it is possible to suppress the second term from (3). By doing this, the complexity of the tie-break strategy has not been modified. Therefore, in the second place, we have to approximate the remaining term: the terms $d_c^k(v_{\sigma(j)})$ are actually determined only after all the nodes have been chosen. On the other hand, the quantities $d_c^{k-1}(v_{\sigma(j)})$ are already at hand. Then there are two obvious ways of approximating (3). The first one is to replace it by

$$f(v_{\sigma(j)}) = \sum_{j=1}^{q} d_c^{k-1}(v_{\sigma(j)}), \tag{4}$$

which amounts to sorting the nodes of the last front by increasing current degree (before they are taken into account) and to keeping the first q. Thus this process can be rewritten as:

Sort the nodes by increasing actualized degree: $(v_{\sigma(j)})_{j=1,\cdots,p_k}$.
Keep $(v_{\sigma(j)})_{j=1,\cdots,q}$. (TB1)

The second way is to choose one node after the other: a node with the lowest current degree is selected, then the current degree of the remaining nodes is updated, a second node is chosen and the process goes on. Thus it can be summarized as:

Sort the nodes by increasing actualized degree: $(v_{\sigma(j)})_{j=1,\cdots,p_k}$ and keep $v_{\sigma(1)}$.
For $l = 2$ to q (TB2)
  Update the current degree of $(v_{\sigma(j)})_{j=1,\cdots,p_k}$, sort and keep $v_{\sigma(l)}$.

It is clear that, in terms of edge cuts, TB1 produces more cuts than TB2 and that, in terms of computational cost, TB2 is more expensive than TB1. As stated in the algorithm describing GP, we have kept the tie-break strategy TB1.

By construction, the heuristic proposed in this paper relies mainly on the iterative construction of the subsets via the fronts, and less on the tie-break strategy. That is the reason why it should be applied mainly to *sparse* graphs, because then it is possible to get many fronts before having to break the tie when enough nodes have been aggregated. On the contrary, for a graph with *high* average degree (or equivalently a *small* diameter), it may happen that the tie-break strategy has to be used right from the start, i.e. without building even a single front but instead having to select some of the nodes of the first front. In this case, using the more expensive tie-break strategy TB2 is advised.

4.3. *Finiteness, effectiveness and complexity*

Let us show that, when GP is applied to the p-edge partitioning problem for a connected graph G, it indeed builds p connected subsets.

Theorem 1
The algorithm GP terminates and always builds p connected subsets.

Proof
First of all, let us prove that the algorithm terminates. Note that for this, we simply have to verify that the value of i is incremented in a finite number of operations (loop 1). Step 1(a) is straightforward. For step 1(b), one has to verify that a node v_s can always be selected. For this, it suffices to check that the current boundary can not be empty: this stems from its definition. If $i = 1$, then it is the original boundary. On the other hand, for $i > 1$, it is clear that there is at least one unmarked node neighboring a marked one: it belongs to the current boundary which is therefore not empty. Finally, one goes through steps 1(c), 1(d) and 1(e) in a finite number of operations (using front techniques and sorting).

In order to prove that it builds exactly p connected subsets, let us proceed by contradiction. Suppose that $\tilde{p}$ connected subsets are constructed, $\tilde{p} \neq p$, when the algorithm terminates. Due to the structure of GP (iterate while $i < p$ and finally build V_p), it is clear that $\tilde{p} < p$. This implies that there exists i, $1 \leq i \leq p$, such that $n_i = 0$.

If $i = p$, then let us prove that we also have $n_{p-1} = 0$ (here, we want to make use of the body of loop 1 and therefore need to prove that $i \leq p - 1$).

As $n_p = 0$ and $V = \bigcup_{i=1}^{n} V_i$,

$$N = \sum_{j=1}^{p-1} n_j. \tag{5}$$

Now, using step 1(a) of the algorithm for $i = p - 1$ gives $n_{p-1} = \lfloor (N - \sum_{j=1}^{p-2} n_j)/2 \rfloor$,

or equivalently

$$2n_{p-1} \le \left(N - \sum_{j=1}^{p-2} n_j \right). \tag{6}$$

This with (5) yields $n_{p-1} = 0$. Thus there exists i, $1 \le i \le p-1$, such that $n_i = 0$. Without changing notations, suppose that i is the smallest index such that $n_i = 0$, i.e. $n_j > 0$, for $1 \le j \le i-1$. Note that V_1 can always be built because G is connected; so $i > 1$.

Then, when $n_i = 0$ is computed ($i < p$), it comes after either step 1(d) for $i - 1$ or step 1(e) for i. Next we shall prove that a contradiction can be reached for these two cases.

In the first case, both the definition of n_{i-1} and the inequality $n_{i-1} \ge 1$ (by assumption on index i) imply

$$[p - (i-2)] \le n_{i-1}[p - (i-2)] \le N - \sum_{j=1}^{i-2} n_j.$$

Therefore,

$$\left[N - \sum_{j=1}^{i-1} n_j \right] [p - (i-2)] = \left[N - \sum_{j=1}^{i-2} n_j \right] [p - (i-2)] - n_{i-1}[p - (i-2)]$$

$$\ge \left[N - \sum_{j=1}^{i-2} n_j \right] [p - (i-1)]$$

$$\ge [p - (i-2)][p - (i-1)], \qquad \text{or}$$

$$N - \sum_{j=1}^{i-1} n_j \ge [p - (i-1)]. \tag{7}$$

Steps 1(a) and (7) yield $n_i \ge 1$, a contradiction.

In the second case, we denote by n_j^{prev} the quantities obtained before step 1(e), i.e. before reassigning the nodes of V_i, and n_j^{curr} the current ones. Let us note that necessarily $n_i^{prev} \ge 2$, because otherwise step 1(e) cannot be reached (if $n_i^{prev} \le 1$, the last step was 1(d)). Thus

$$2[p - (i-1)] \le n_i^{prev}[p - (i-1)] \le N - \sum_{j=1}^{i-1} n_j^{prev}.$$

Finally, denote by n_i^{conn} the number of reassigned nodes. Then

$$n_i^{conn} < n_i^{prev} \quad \text{and} \quad \sum_{j=1}^{i-1} n_j^{prev} + n_i^{conn} = \sum_{j=1}^{i-1} n_j^{curr}.$$

Therefore

$$\left[N - \sum_{j=1}^{i-1} n_j^{curr}\right][p - (i-1)] = \left[N - \sum_{j=1}^{i-1} n_j^{prev} - n_i^{conn}\right][p - (i-1)]$$

$$\geq \left[N - \sum_{j=1}^{i-1} n_j^{prev} - n_i^{prev}\right][p - (i-1)]$$

$$\geq \left[N - \sum_{j=1}^{i-1} n_j^{prev}\right][p - i]$$

$$\geq 2[p - (i-1)][p - i], \qquad \text{implying}$$

$$N - \sum_{j=1}^{i-1} n_j^{curr} \geq [p - (i-1)]. \tag{8}$$

Steps 1(a) and (8) yield $n_i^{curr} \geq 1$, again a contradiction. $\qquad\square$

Let us prove now that GP is inexpensive.

Theorem 2
The complexity of GP is $O(M)$.

Proof
In order to obtain this result, let us consider the complexity of the algorithm step by step for the construction of a given subset V_i, $i < p$.

The complexity of step 1(a) is $O(1)$.

For step 1(b), we shall consider two cases, depending on whether the current boundary is new or not.

If the jth current boundary is new, then v_s is chosen in $O(n_b^j)$ operations, using 1(b)i. and 1(b)iii. Note that a node belongs to at most one current boundary, so $\sum_j n_b^j \leq N$. Moreover, if the current boundary is multiconnected, then the choice is made in $O(n_b^{j,l})$ operations for the lth connected component, with $\sum_l n_b^{j,l} = n_b^j$.

If the current boundary is not new, then let us divide the operations count into two subcases, according to whether it comes after step 1(d) for $i - 1$ or step 1(e) for i. In the first subcase, it is a simple matter to check that v_s can be chosen in $O(\sum_{v \in V_{i-1}} d_c(v))$ operations, by using 1(b)ii and then 1(b)i and 1(b)iii.

If the second subcase happens, then necessarily it comes after a first subcase for which we suppose that a sorted list of candidates has been built. This can be done in $O(\sum_{v \in V_{i-1}} d_c(v))$ operations, if one recalls that the nodes are sorted by current degree. Indeed, the current degree of a node varies between 0 and d_{max}, therefore sorting n of them is done in $O(n)$ operations by piling them up in $(d_{max} + 1)$

heaps. Thus, to find v_s, it suffices to update this sorted list, which can be done in $O(\sum_{v \in V_i^{prev}} d_c(v))$ operations.[3]

By dividing the case in this way, note that a given node appears at most twice in the summation (once while examining its unmarked neighbors, and once while building the sorted list).

For step 1(c), it is clear that it is done in $O(m_i)$ operations (construction by fronts).

In step 1(d), updating the current and virtual boundaries is done respectively in $O(n_i)$ and $O(m_i)$ operations and the tie-break strategy requires $O(n_i)$ operations, if one again takes the sorting in heaps into account.

Finally, if the occasion should arise (step 1(e)), nodes are reassigned in $O(m_i^{prev})$ operations by successively reassigning each of them (in reverse order) to the subset which holds the highest number of its neighbors.

Clearly, the complexity of step 2 is $O(m_p)$.

To recapitulate, the complexities of the p steps 1(a), 1(b), 1(c), 1(d), 1(e) and of step 2 are respectively bounded by:

$$O(p), \qquad O\left(\sum_j n_b^j + O\left(\sum_i \sum_{v \in V_i} d_c(v)\right)\right),$$

$$O\left(\sum_i m_i\right), \qquad O\left(\sum_i n_i + \sum_i m_i\right),$$

$$O\left(\sum_i m_i\right), \qquad O(m_p).$$

The conclusion follows if one uses $O(\sum_i \sum_{v \in V_i} d_c(v)) \leq O(\sum_i m_i) \leq O(M)$. $\square$

4.4. *Case of a multiconnected graph*

In the case of a multiconnected graph G, an $O(M)$ algorithm (breadth-first search, fronts, etc.) is used to find its maximal connected components $(G_i)_{i=1,\cdots,q}$. Then GP is applied on each component with the proper number of subsets (proportional to the number of nodes of the component). Thus for a multiconnected graph the algorithm is of the following form.

Step 1. Compute the maximal connected components $(G_i)_{i=1,\cdots,q}$ of G.
Step 2. Compute $(p_i)_{i=1,\cdots,q}$ such that p_i is proportional to N_i and $\sum_{i=1}^q p_i = p$.
Step 3. For $i = 1$ to q, partition G_i with GP.

---

[3] If this list is empty, a new starting node v_c, which belongs to the jth current boundary, is selected. The cost of selecting (possibly several times) v_c is $O(n_b^j)$. Indeed, initially, a list of the nodes of this current boundary is built. Then, whenever necessary, v_c can be defined as the first node in the list which still belongs to the current boundary.

Clearly, theorem 1 can be extended to the multiconnected case, as it holds for each of the connected components G_i. Theorem 2 can also be applied in this case, because the complexity of finding the maximal connected components is $O(M)$.

5. Postprocessing steps

As previously mentioned, the added connectivity often produces unbalanced subsets, in terms of (n_i). Thus there is a need for postprocessing, the goal of which is to balance the subsets, i.e. decrease the value of $\sigma = \sqrt{(1/p)\sum_{i=1}^{p}(n_i - \bar{n})^2}$. In this respect, it differs from more classical post-treatments which come after well balanced subsets have been obtained: in that case, the emphasis is on reducing the number of edge cuts m_0. Hereafter, we briefly recall some of these classical methods and then consider our balancing heuristics. All these postprocessing steps are based on node interchanging.

5.1. *Reducing the number of edge-cuts*

To our knowledge, the first heuristic was proposed by Kernighan and Lin [24]. In that case, the quality of a *balanced* 2-edge partition $V_1 \cup V_2$ is improved by exchanging subsets of V_1 and V_2. The selection criterion of the subsets is determined from a gain function which is defined as follows. First, for $v \in V_1$, $g_v = d_{V_2}(v) - d_{V_1}(v)$ 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:

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

So, after computing for each node $v \in V_1$ and $w \in V_2$ the values of g_v and g_w, the algorithm chooses a pair of nodes (v_1, w_1) which maximizes the gain g_{v_1,w_1}. Then, for all neighbors u of v_1 or w_1 the values of g_u are updated. The optimizing process is iterated until $g_{v_{n-1},w_{n-1}}$ is computed. Finally, the algorithm chooses among all pairs of subsets $\{v_1, v_2, \ldots, v_k\}$, $\{w_1, w_2, \ldots, w_k\}$ for $k \in \{1, \ldots, n-1\}$ the one whose swapping corresponds to the maximum of the gain $\sum_{l=1}^{k} g_{v_l,w_l}$.

This process can be repeated iteratively until no better improvement, regarding the number of intersubset edges, is obtained. The complexity of one iteration is $O(N^2 \log N)$. The number of iterations is not bounded a priori.

Many variants of this heuristic have been proposed in the past years which aim at reducing the overall complexity. Fiduccia and Mattheyses [15] remarked that the high cost of Kernighan and Lin's heuristic is due to the $\delta(v, w)$ term in the expression of $g_{v,w}$. Thus, instead of selecting nodes by pairs, they proposed to choose them

one after the other, that is v_1, then w_1, then v_2 etc. Then the gain is simply:

$$g'_{v,w} = g_v + g_w,$$

where g_w takes into account the fact that v has already been moved from V_1 to V_2. Again, the process can be repeated. The complexity of one iteration is reduced to $O(M)$. Still, the number of iterations is not bounded.

Recently, Hendrickson and Leland [22,23] generalized Fiduccia and Mattheyses' method to an arbitrary number of subsets. Instead of a single gain function there are $(p - 1)$ gain functions associated with each node. Each of these functions computes the gain obtained by moving the given node to a specific subset. Moreover, their method integrates other features. First, their gain function takes into consideration an intersubset cost metric. Second, their starting partition is not necessarily balanced, so they give precedence to moves from large subsets to small ones. The complexity of one iteration is $O((p - 1)M)$, with no bound for the number of iterations.

Also, Dutt [12] improved Kernighan and Lin's heuristic in the case of random sparse graphs. In this case, he proved that the search for optimal pairs can be done at low cost if the data structure has been built carefully. The complexity of one iteration is $O(M \log(N))$ but reduces to $O(\bar{d}M)$ for graphs such that the average degree $\bar{d}$ is less than $O(\log(N))$. Again, there is no bound for the number of iterations.

5.2. Balancing and reshaping the subsets

We now present our postprocessing heuristic which we named a *retrofitting* method after the term used by engineers when bracing existing structures. This method was first presented in [10] and has been naturally suggested by numerous numerical experiments conducted with GP. This heuristic is made of three steps. Steps 1 and 3 are identical *reshaping* steps while step 2 is a *balancing* step.

The reshaping step tries to redesign the outlines of subsets by deleting the *excrescences*, i.e. nodes that are attached to the subset to which they belong by a single edge. From a geometric point of view, 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. Moreover, the value of m_0 can be improved but cannot deteriorate. Indeed, the number of edges that are cut from a vertex v goes from $d(v) - 1$ (by definition of an excrescence) to $d(v) - d_{V_k}(v)$, where V_k is the subset to which v is reassigned, and $d_{V_k}(v) \geq 1$.

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 (in practice five iterations). Here, one iteration consists of spanning the whole set of nodes V and reassigning the excrescences.

As we previously remarked, 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 balancing the subsets by moving nodes from large subsets to small ones. One iteration of this consists of three parts (a), (b) and (c). First, in part (a), the largest and smallest subsets are determined. Then nodes of the largest subset are reassigned to its smallest neighbor in part (b). Last, in part (c), nodes are reassigned to the smallest subset, coming from its largest neighbor.

For this balancing step we have two strategies. On the one hand, the goal of the first method LB1, presented in [10], is to decrease the value of σ, i.e. to balance the subsets, while trying to reduce m_0, the number of edge cuts. Let us detail parts (b) and (c) for LB1. For this method, a *single* node is chosen and moved to reduce m_0. Let us call V_i the subset from which a node v is taken and V_j the subset to which it is reassigned. Then v is selected among the nodes of V_i bordering V_j. It is the one which leads to the best partition in terms of edge cuts. In other words, it is a node v of V_i which satisfies $d_{V_j}(v) \neq 0$ and maximizes $\{d_{V_j}(v) - d_{V_i}(v)\}$. Moreover, before moving a node v from subset V_i one has to make sure that $V_i \backslash \{v\}$ remains connected. For that, proposition 3 is used. If this is not the case, then another boundary node of V_i is selected. If no node of V_i bordering V_j satisfies this connectivity requirement, then another subset (V_j in part (b), V_i in part (c)) has to be chosen.

On the other hand, for the second method LB2, the selection of the nodes that are to be reassigned in parts (b) and (c) differs from LB1 in that a whole *front* of nodes is moved (following the main idea used for the GP heuristic). Thus, instead of reassigning a single boundary node, all boundary nodes, i.e. all nodes v of V_i such that $d_{V_j}(v) \neq 0$, are moved. So, balancing prevails over minimizing the number of edge cuts. This set of boundary nodes corresponds exactly to $F(V_j, 1) \cap V_i$. Moreover, as in the previous method, one has to check that $V_i \backslash F(V_j, 1)$ is connected. This is done using proposition 3. If this is not the case, then another subset has to be selected. Evidently, this process is valid only when the standard deviation $\sigma = \sqrt{(1/p) \sum_{i=1}^{p} (n_i - \bar{n})^2}$ decreases. If we use a prime ($'$) for new quantities, call δn the number of nodes of the front and let $\delta \sigma = p(\sigma'^2 - \sigma^2)$, we have

$$\delta \sigma = (n_i' - \bar{n})^2 + (n_j' - \bar{n})^2 - (n_i - \bar{n})^2 - (n_j - \bar{n})^2$$

$$= (n_i - \bar{n} - \delta n)^2 + (n_j - \bar{n} + \delta n)^2 - (n_i - \bar{n})^2 - (n_j - \bar{n})^2$$

$$= 2\delta n\{\delta n - (n_i - n_j)\}.$$

In other words, LB2 is used when $\delta n \leq n_i - n_j$.

For both balancing steps, the process is iterated until the standard deviation approaches the optimal standard deviation

$$\sigma_{opt} = \sqrt{\frac{1}{p}[((\bar{n}+1)p - N)(\bar{n} - x)^2 + (N - \bar{n}p)((\bar{n}+1) - x)^2]},$$

where $x = N/p$, or does not decrease anymore (in practice after five iterations).

The complexity of one iteration of the reshaping step is $O(M)$. The complexity of one iteration of the balancing step (either LB1 or LB2) is $O(p)$ for part (a) and $O(m_i + m_j)$ for parts (b) and (c). Note that, as for the postprocessing steps described in subsection 5.1, when the retrofitting method is performed, neither the number of iterations for the reshaping steps nor for the balancing step are bounded a priori. Nevertheless, it is expected that the number of balancing steps which are performed when using LB2 is much smaller than when using LB1.

6. Examples

Many numerical experiments have already been presented by the authors on physical meshes for the GP heuristic on the one hand [8,9] and GP with the retrofitting method using LB1 on the other hand [10]. In [10], a number of comparisons were made with spectrally derived heuristics. However, on a few tests the balancing time was important. This led to the design of a new balancing step (LB2), leading to a heuristic totally based on fronts, which it was hoped would be faster.

The meshes for which we compare the greedy heuristics have the following characteristics:

Annulus	$N = 8448,$	$M = 33024.$
Airfoil	$N = 258990,$	$M = 773168.$
Spheres	$N = 9020,$	$M = 59418.$
Cube	$N = 15625,$	$M = 186696.$
Eppstein	$N = 131137,$	$M = 392256.$

Annulus, Airfoil and Eppstein are 2-dimensional meshes, whereas Spheres and Cube are 3-dimensional. For more details, we refer the reader to [10]. Hereafter we list in 5 tables the result of the comparison of GP alone or with the retrofitting method using either LB1 or LB2 (called respectively GP_1 and GP_2). The parameters listed below are the *average standard deviation* σ/n (%), the percentage of edge cuts m_0/M and the elapsed time t in seconds (on a 50 MHz TI "Viking" SuperSparc processor).

Table 1
Annulus.

p	σ/n (%)			m_0/M (%)			t (s)		
	GP	GP_1	GP_2	GP	GP_1	GP_2	GP	GP_1	GP_2
4	5.8	0.3	0.3	2.69	2.63	2.64	0.08	0.52	0.35
16	0	0	0	5.86	5.83	5.83	0.08	0.19	0.18
64	11.2	4.4	4.1	11.87	11.72	11.85	0.08	0.72	0.43
256	6.3	2.9	2.9	25.77	24.75	24.79	0.08	0.49	0.38

Table 2
Airfoil.

p	σ/n (%)			m_0/M (%)			t (s)		
	GP	GP_1	GP_2	GP	GP_1	GP_2	GP	GP_1	GP_2
16	9.8	0	0	0.95	1.00	1.00	3.0	300.0	5.7
64	17.2	6.7	6.7	1.87	1.88	1.92	2.9	46.0	8.7
256	8.8	4.3	3.0	4.07	4.08	4.11	2.8	17.9	14.0
1024	5.2	3.0	2.1	8.45	8.35	8.37	3.0	16.2	12.8

Table 3
Spheres.

p	σ/n (%)			m_0/M (%)			t (s)		
	GP	GP_1	GP_2	GP	GP_1	GP_2	GP	GP_1	GP_2
4	0	0	0	8.03	8.03	8.03	0.18	0.26	0.25
16	0.1	0	0.1	18.37	18.30	18.29	0.18	0.30	0.30
64	6.7	0.8	0.4	29.91	29.81	29.83	0.18	0.78	0.50
256	3.0	1.8	1.6	48.24	47.51	47.54	0.19	0.44	0.41

Table 4
Cube.

p	σ/n (%)			m_0/M (%)			t (s)		
	GP	GP_1	GP_2	GP	GP_1	GP_2	GP	GP_1	GP_2
4	0	0	0	7.43	7.43	7.43	0.41	0.66	0.62
16	3.9	0.1	0.1	16.36	16.35	16.37	0.45	2.8	2.4
64	0.3	0.2	0.2	27.26	27.22	27.22	0.42	0.76	0.76
256	0.9	0.5	0.5	46.13	46.06	46.05	0.45	1.0	0.92

Table 5
Eppstein.

p	σ/n (%)			m_0/M (%)			t (s)		
	GP	GP_1	GP_2	GP	GP_1	GP_2	GP	GP_1	GP_2
16	3.3	0	0	1.35	1.35	1.35	1.2	7.7	3.7
64	5.7	3.5	2.1	2.96	2.96	2.96	1.1	6.2	5.1
256	1.8	1.1	1.1	6.07	6.01	6.00	1.1	3.9	3.4
1024	2.2	1.4	1.4	12.30	12.03	12.02	1.2	4.9	4.3

As expected, when GP is used alone, the elapsed time is independent of p. For low average degree graphs (2D meshes), it generally produces some imbalance between subsets. Moreover, the retrofitting method is more expensive in these cases.

We can make several other remarks, when the retrofitting method is used. First, the average load balancing is improved when using LB2 over LB1. Second, the number of edge cuts m_0 is only marginally higher for LB2, although it does not take into account the optimization of m_0. Third, elapsed times are always smaller and, in addition, LB2 produces in some cases results equivalent to those obtained with LB1 much faster. Also, when LB2 is used for a given mesh, there are no large fluctuations in terms of elapsed time for different values of the number of subsets (see the Airfoil). Finally, let us mention that if we divide the number of edges by the average of the elapsed time (for a given mesh and different values of p), then we have, for the partitioning heuristic GP_2, the following results[4]:

Annulus $M/t \approx 10 \times 10^4$ edges/s.

Airfoil $M/t \approx 8 \times 10^4$ edges/s.

Spheres $M/t \approx 16 \times 10^4$ edges/s.

Cube $M/t \approx 16 \times 10^4$ edges/s.

Eppstein $M/t \approx 10 \times 10^4$ edges/s.

There are two behaviors, one for low average degree graphs (Annulus, Airfoil, Eppstein), and the other for higher average degree graphs (Spheres, Cube). As previously noted, the retrofitting is less expensive in the latter case because GP provides better balanced partitions.

7. Conclusions

In this paper, we have presented a heuristic which approximates the p-edge partitioning problem for large sparse graphs G with a connectivity constraint on the subsets. It is decomposed into two parts: the computation of an initial partition and an iterative retrofitting process. We proved the finiteness and correctness of the initial step and also showed that its complexity is $O(M)$, where M is the number of edges of the graph, by using front-based tools. For the retrofitting method, we designed a new balancing step, also based on front techniques, in order to reduce the number of iterations of the method. Numerically, we verified on various examples that the new retrofitting process leads to reduced elapsed time to compute the partition, and that these elapsed times are linear with respect to M.

---

[4] For GP alone, M/t is in the order of 36×10^4 edges/s.

References

[1] C.J. Alpert and A.B. Kahng, Multi-way partitioning via spacefilling curves and dynamic programming, in: *31st Design Automation Conference* (ACM/IEEE, 1994) pp. 652–657.

[2] 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-092-033 (1992).

[3] E.R. Barnes, An algorithm for partitioning the nodes of a graph, SIAM J. Alg. Discr. Meth. 3 (1982) 541–550.

[4] E.R. Barnes, A. Vannelli and J.Q. Walker, A new heuristic for partitioning the nodes of a graph, SIAM J. Discr. Meth. 1 (1988) 299–305.

[5] C. Berge, *Graphes* (Gauthier-Villars, Paris, 1983).

[6] T.N. Bui, S. Chaudhuri, F.T. Leighton and M. Sipser, Graph bisection algorithm with good average case behavior, Combinatorica (1987) 171–191.

[7] N. Chistofides and P. Brooker, The optimal partitioning of graphs, SIAM J. Appl. Math. 30 (1976) 55–69.

[8] P. Ciarlet, Jr. and F. Lamour, Un algorithme rapide de partitionnement automatique de graphes, Technical report, CEA N-2738 (1993).

[9] P. Ciarlet, Jr. and F. Lamour, An efficient low cost greedy graph partitioning heuristic, Technical report, UCLA, CAM 94-1 (1994).

[10] 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).

[11] P. Ciarlet, Jr., F. Lamour and B.F. Smith, On the influence of the partitioning schemes on the efficiency of overlapping domain decomposition methods, in: *5th Symp. on the Frontiers of Massively Parallel Computation* (IEEE, 1995) pp. 375–384.

[12] S. Dutt, New faster Kernighan-Lin graph partitioning algorithms, in: *Int. Conf. on CAD-93* (IEEE/ACM, 1993) pp. 370–377.

[13] C. Farhat, A simple and efficient automatic FEM domain decomposer, Comp. Structures 28 (1988) 579–602.

[14] 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).

[15] C.M. Fiduccia and R.M. Mattheyses, A linear-time heuristic for improving network partitions, in: *19th Design Automation Conf.* (IEEE, 1982) pp. 175–181.

[16] M. Fiedler, Algebraic connectivity of graphs, Czech. Math. J. 23 (1973) 298–305.

[17] L.R. Foulds, *Graph Theory Applications* (Springer, 1992).

[18] A. George and J.W.H. Liu, An automatic nested dissection algorithm for irregular finite element problems, SIAM J. Numer. Anal. 15 (1978) 1053–1069.

[19] L. Hagen and A.B. Kahng, A new approach to effective circuit clustering, in: *IEEE/ACM Int. Conf. on Computer-Aided Design* (IEEE, 1992) pp. 422–427.

[20] L. Hagen and A.B. Kahng, New spectral methods for ratio cut partitioning and clustering, IEEE Trans. Comp. Aided Design 11 (1992) 1074–1085.

[21] B. Hendrickson and R. Leland, An improved spectral graph algorithm for mapping parallel computations, Technical report, Sandia, SAND 92-1460 (1992).

[22] B. Hendrickson and R. Leland, Multidimensional spectral load balancing, Technical report, Sandia, SAND 93-0074 (1993).

[23] B. Hendrickson and R. Leland, A multilevel algorithm for partitioning graphs, Technical report, Sandia, SAND 93-1301 (1993).

[24] B.W. Kernighan and S. Liu, An efficient heuristic for partitioning graphs, Bell Syst. Techn. 49 (1970) 291–307.

[25] J.W.H. Liu, A graph partitioning algorithm by node separators, ACM Trans. Math. Softw. 15 (1989) 198–219.

[26] G.L. Miller, S-H. Teng, W. Thurston and S.A. Vavasis, Automatic mesh partitioning, in: *Graph Theory and Sparse Matrix Computations* (Springer, 1993) pp. 57–84.

[27] A. Pothen, H.D. Simon and K.P. Liou, Partitioning sparse matrices with eigenvectors of graphs, SIAM J. Matrix Anal. Appl. 11 (1990) 430–452.

[28] F. Rendl and H. Wolkowicz, A projection technique for partitioning the nodes of a graph, Technical report, University of Waterloo, CORR 90–20 (1990).

[29] C. Roucairol and P. Hansen, Problème de la bipartition minimale d'un graphe, RAIRO 21 (1987) 325–348.

[30] H.D. Simon, Partitioning of unstructured problems for parallel processing, Comp. Syst. Eng. 2 (1991) 135–148.