

Level-based heuristics and hill climbing for the antibandwidth maximization problem

Jennifer Scott¹ and Yifan Hu²

ABSTRACT

The antibandwidth maximization problem aims to maximize the minimum distance of entries of a sparse symmetric matrix from the diagonal and as such may be regarded as the dual of the well-known bandwidth minimization problem. In this paper, we consider the feasibility of adapting heuristic algorithms for the bandwidth minimization problem to the antibandwidth maximization problem. In particular, using an inexpensive level-based heuristic we obtain an initial ordering that we refine using a hill-climbing algorithm. This approach performs well on matrices coming from a range of practical problems with an underlying mesh. Comparisons with existing approaches show that, on this class of problems, our algorithm can be competitive with recently reported results in terms of quality while being significantly faster and applicable to much larger problems.

Keywords: antibandwidth maximization, sparse matrices, level sets, hill climbing.

AMS(MOS) subject classifications: 65F30, 65F50

¹ Computational Science and Engineering Department, Rutherford Appleton Laboratory, Chilton, Oxfordshire, OX11 0QX, UK.

Email: jennifer.scott@stfc.ac.uk

Work supported by EPSRC grant EP/E053351/1.

² AT&T Labs Research, 180 Park Avenue, Florham Park, New Jersey, NJ 07932, USA.

Email: yifanhu@research.att.com

1 Background and motivation

Since the 1960s, considerable attention has been paid to the design and development of algorithms for minimizing the bandwidth of a sparse symmetric matrix $A = \{a_{ij}\}$. That is, finding a labelling (or ordering) of the rows and columns of A that minimizes the maximum distance b from the diagonal

$$b = \min_i \{ \max_j \{ |i - j| : a_{ij} \neq 0 \} \}$$

(see, for example, [2], [6], [11], [12], [14]). Until relatively recently, much less attention has focused on the antibandwidth maximization problem, which is the problem of finding a labelling of the rows and columns of A that maximizes the minimum distance ab from the diagonal

$$ab = \max_i \{ \min_j \{ |i - j| : i \neq j \text{ and } a_{ij} \neq 0 \} \}.$$

Many algorithms for reducing the bandwidth of A make extensive use of the adjacency graph \mathcal{G} of A . This is an undirected graph that has a node for each row (or column) of the matrix and node i is a neighbour of node j if a_{ij} (and by symmetry a_{ji}) is an entry (nonzero) of A . In terms of graphs, the antibandwidth maximization problem is to label the nodes of the graph such that the length of the shortest edge is maximized (that is, the labelling difference of the end nodes among all edges is maximized). This problem was introduced by Leung et al. [10] in 1984 in connection with multiprocessor scheduling problems. It is also referred to as the *dual bandwidth* problem [20] or the *separation* problem [10]. It arises in a number of practical applications. For example, it belongs to the family of obnoxious facility location problems [1]. Here the “enemy” graph is represented by n people and there is an edge between two people iff they are enemies. The problem is to build each person a house along a road so that the minimal distance between enemies is maximized. Another example is the radio frequency assignment problem in which the nodes correspond to transmitters and the edges are between interfering transmitters; the objective is to assign the frequencies so that those for the interfering transmitters are as different as possible.

Like the bandwidth minimization problem, the antibandwidth maximization problem is NP-Complete [10]. In the literature, theoretical results have been presented for some special graphs, including paths, cycles, rectangular grids, special trees and complete bipartite graphs (see, for example, [16] and the references therein). Recently, there has been an interest in developing algorithms to compute solutions that are close to the optimal and that apply to general graphs. In particular, Hu, Kobourov and Veeramoni [7] have developed an algorithm **GSpectral** (Greedy Spectral) that is based on computing the eigenvector corresponding to the largest eigenvalue of the Laplacian associated with the graph and then using a greedy refinement algorithm. They have applied this to the maximum differential graph colouring problem and reported results for some small examples. Duarte, Martí, Resende and Silva [4] have proposed an integer linear programming formulation and several heuristics based on **GRASP** (Greedy Randomized Adaptive Search Procedure) with path relinking. They present some high-quality computational results for general graphs, although the run-times for their relatively modest-sized test problems (graphs with fewer than 9000 nodes) are quite high (typically several minutes for their fastest approach applied to their largest problems). Thus we would like to develop alternative algorithms for increasing the antibandwidth that are significantly faster while retaining good quality. This paper is the first step in achieving this aim.

An important and well-known example of a bandwidth reduction algorithm is the Cuthill-McKee algorithm [2] and its many variants, including the Gibbs-Poole-Stockmeyer algorithm [6]. The Cuthill-McKee algorithm constructs a level set structure of the graph \mathcal{G} and labels the nodes according to these levels. In this paper, we consider the feasibility of modifying this approach to obtain a practical algorithm for increasing the antibandwidth of A . We find that on its own this is not generally sufficient to yield large antibandwidths but that, when combined with a suitable refinement algorithm, we are able to compute high-quality orderings for problems that arise from a range of applications with an underlying mesh. Furthermore, our approach is fast and thus potentially practical for larger problems than cannot be tackled by either the **GSpectral** or **GRASP** approaches.

The rest of this paper is organised as follows. We begin (Section 2) by briefly recalling the Cuthill-McKee algorithm and then consider how it might be modified for the antibandwidth maximization problem. In Section 3, we look at modifying the hill-climbing algorithm of Lim, Rodrigues and Xiao [11] to improve a given ordering. In Section 4, our proposed algorithms are used to reorder a set of test matrices and our results are compared with those of for `GSPectral` and `GRASP` in Section 5. We summarise our findings and discuss future work in Section 6.

2 Level-based approach to antibandwidth problem

In this section, we first recall the Cuthill-McKee algorithm for bandwidth minimization and then consider how it may be adapted for the antibandwidth maximization problem.

2.1 The Cuthill-McKee algorithm

Given a starting node s , the Cuthill-McKee algorithm proceeds by relabelling the nodes of the adjacency graph \mathcal{G} by order of increasing distance from s . The algorithm is outlined in Figure 2.1. Here the degree of a node i is defined as the number of neighbours it has (that is, the number of nodes $j \neq i$ for which $a_{ij} \neq 0$). If \mathcal{G} has more than one component, the procedure is repeated from a starting node in each component.

```

Algorithm 1: Cuthill-McKee
Label  $s$  as node 1;  $l_1 = \{s\}$ ;  $i = 1$ 
do  $k = 2, 3, \dots$  until  $i = n$ 
     $l_k = \{\}$ 
    do for each  $v \in l_{k-1}$  in label order
        do for each neighbour  $u$  of  $v$  that has not been labelled,
            in order of increasing degree
                add  $u$  to  $l_k$ ;  $i = i + 1$ ; label  $u$  as node  $i$ 
        end do
    end do
end do

```

Figure 2.1: Cuthill-McKee ordering algorithm

Ordering the nodes in this way groups them into *level sets*, that is, nodes at the same distance from the starting node. Since nodes in level set l_k can have neighbours only in level sets l_{k-1} , l_k , and l_{k+1} , the reordered matrix is block tridiagonal with blocks corresponding to the level sets. It is therefore desirable that the level sets be small, which is likely if there are many of them. The level-set structure rooted at s is denoted by $L(s) = \{l_1, l_2, \dots, l_h\}$. Algorithms for finding a good starting node are usually based on finding a pseudo-diameter of \mathcal{G} (a pair of nodes that are a maximum distance apart or nearly so). Much effort has gone into efficiently finding a pseudo-diameter; algorithms are generally based either on using level sets (see, for example, [6] and [17] and the references therein) or using the Fiedler vector [15]. The (non-unique) optimal bandwidth minimization ordering that is obtained by applying the Cuthill-McKee algorithm to a 5×5 mesh is illustrated in Figure 2.2. Here the letters indicate the initial ordering and the numbers in parentheses are the new labels. A grey scale colouring is used, with the first node given the lightest shade and the last node the darkest shade. Node ‘a’ is chosen as the starting node (but any of the other corner nodes could equally well have been chosen).

2.2 Level-based approach for antibandwidth maximization problem

In their paper, Miller and Pritikin [13] establish tight bounds for the maximum antibandwidth for various classes of graphs. Further bounds were recently presented by Raspaud et al. [16]. In particular, Raspaud

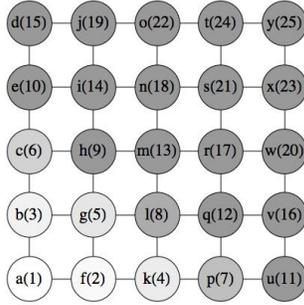


Figure 2.2: Cuthill-McKee labelling of 5×5 mesh. Here the letters indicate the initial ordering and the numbers in parentheses are the new labels. A grey scale colouring is used, with the first node given the lightest shade and the last node the darkest shade.

et al. showed that for a 2-dimensional $m \times k$ mesh with $m \geq k \geq 2$ the lower bound on the maximum antibandwidth proved by Miller and Pritikin is precise, that is, the maximum antibandwidth for such problems is

$$ab = \left\lceil \frac{k(m-1)}{2} \right\rceil. \quad (2.1)$$

Miller and Pritikin describe how this bound can be achieved. Nodes i and j with coordinates (p, q) and (p', q') in the mesh are neighbours if and only if $p = p'$ and $|q - q'| = 1$ or $q = q'$ and $|p - p'| = 1$. Miller and Pritikin set the origin $(0, 0)$ to a corner of the mesh and define $X = \{(p, q) : p + q \text{ is odd}\}$ and $Y = \{(p, q) : p + q \text{ is even}\}$. X is ordered lexicographically and then Y is ordered lexicographically. This is equivalent to choosing the starting node s to be a corner of the mesh, constructing the level-set structure $L(s)$, and then taking the level sets in the order $\{l_2, l_4, \dots, l_h, l_1, l_3, \dots, l_{h-1}\}$ (here h is assumed to be even) and ordering nodes in each level set in turn, in natural order.

For a 3-dimensional $m \times m \times m$ mesh Török and Vrt'o [19] show that

$$ab = \frac{4m^3 - 3m^2}{8} + \mathcal{O}(m). \quad (2.2)$$

Their algorithm for labelling the nodes to achieve this optimal value (up to the third order term) again labels the even numbered level sets and then the odd numbered level sets, starting from a corner.

Choosing to start at a corner of the mesh is equivalent to selecting the starting node to be an end point of a diameter of \mathcal{G} . This suggests that for more general problems we should select s to be an end point of a pseudodiameter, construct $L(s)$ and use the level sets to guide the relabelling. In the Cuthill-McKee algorithm, at each stage the list of candidate nodes for the next label comprises the unlabelled neighbours of the nodes that have already been labelled. Thus a node and its neighbours receive labels that are close to each other, yielding a narrow bandwidth. When attempting to increase the antibandwidth, we need to do the opposite, that is, if a node has been labelled, avoid labelling its neighbours for as long as possible. The approach of Miller and Pritikin does exactly that for mesh problems because none of the nodes in each of the level sets l_r has neighbours in the same level set (the neighbours all belong to the level sets l_{r-1} and l_{r+1}). For more general problems, we will have neighbours belonging to the same level set and so we need to use a strategy to avoid labelling these neighbours too soon. Our algorithm is outlined in Figure 2.3.

Algorithm 2 has a number of sweeps, or passes, through the level-set structure. On each sweep, a flag is used to indicate whether or not a node that has yet to be labelled is a candidate for labelling during that sweep. Initially, all flags are set to zero. When a node u is labelled, all the neighbours v of u that are unlabelled are flagged with the current sweep number; these nodes are not candidates for labelling until

Algorithm 2: Level-based antibandwidth ordering (LB)Input: starting node s , rooted level-set structure $L(s) = \{l_1, l_2, \dots, l_h\}$.Initialise: $sweep = 0$; $flag(1 : n) = 0$; $i = 0$.

```

do until  $i = n$ 
   $sweep = sweep + 1$ 
  do  $r = 1, \dots, h$ 
    do for each unlabelled  $u \in l_r$ 
      if ( $flag(u) = sweep$ ) cycle
         $i = i + 1$ ; give node  $u$  label  $i$ 
        Set  $flag(v) = sweep$  for each unlabelled neighbour  $v$  of  $u$ 
      end do
    end do
  end do
end do

```

Figure 2.3: Level-based antibandwidth ordering algorithm

the next sweep. Note that nodes in level set l_r can only have neighbours in l_{r-1} , l_r and l_{r+1} . For the mesh problems of Miller and Pritikin, Algorithm 2 reduces to ordering the even numbered level sets and then the odd numbered level sets. The run time complexity of Algorithm 2 is $\mathcal{O}(d_{max} * nz)$, where d_{max} is the maximum degree of a node and nz is the number of edges.

The use of Algorithm 2 to order a 5×5 mesh is illustrated in Figure 2.4. Again, the letters indicate the initial ordering and the numbers in parentheses are the new labels; node ‘a’ is chosen as the starting node.

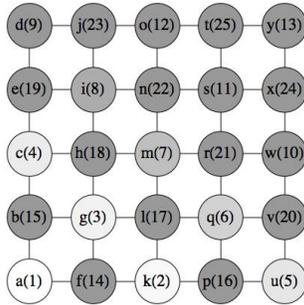


Figure 2.4: LB labelling of 5×5 mesh. Here the letters indicate the initial ordering and the numbers in parentheses are the new labels. A grey scale colouring is used, with the first node given the lightest shade and the last node the darkest shade.

The success of Algorithm 2 depends on the choice of the starting node s and the sizes of the level sets. The rationale for choosing s to be an endpoint of a pseudodiameter is that it will tend to lead to a long thin level-set structure that is hopefully also well-balanced in the sense that (with the exception of the first and last few levels) the levels each have a similar number of entries. We want to avoid $L(s)$ having one level set (or a small number of level sets) l_r that is larger than the other level sets. For suppose a node $u \in l_r$ has a large number of neighbours, most of which also belong to l_r . Once u has been labelled, if there remains no unlabelled nodes in the other level sets, the remaining unlabelled nodes in l_r will be labelled consecutively. This results in an antibandwidth of 1, even though with the new labelling

$$\min_j \{|i - j| : i \neq j \text{ and } a_{ij} \neq 0\}$$

is significantly larger than 1 for all $i \neq n - k$ for small k . This is illustrated in Section 4.1. To help assess

the quality of our labelling, we define the *average antibandwidth* as

$$av = \frac{1}{n} \sum_i \min_j \{|i - j| : i \neq j \text{ and } a_{ij} \neq 0\}. \quad (2.3)$$

Note this definition considers upper and lower triangular entries and thus maximizing av is **not** analogous to the problem of minimizing the profile of A . Recall that the profile is defined to be

$$profile = \frac{1}{n} \sum_i \max_j \{i - j : j < i \text{ and } a_{ij} \neq 0\} \quad (2.4)$$

and thus involves minimizing the sum of the maximum distance from the diagonal in the lower triangular part of A only.

3 Hill climbing refinement algorithm

The level-based algorithm is a global algorithm. In this section, we look at how we might locally refine the ordering to improve the antibandwidth. Lim et al [11] propose a hill-climbing algorithm for reducing the bandwidth of a symmetric matrix. This local search strategy was adapted for unsymmetric matrices by Reid and Scott [18]. Here we propose using hill climbing to increase the antibandwidth.

The idea behind hill climbing is that at each step a search is made for a non-critical node to swap with a critical node. For the antibandwidth maximization problem, i is defined to be *critical* if

$$\min_k \{|i - k| : i \neq k \text{ and } a_{ik} \neq 0\} = ab. \quad (3.1)$$

If i is critical, we look for a non-critical j such that symmetrically permuting i and j (that is, swapping rows i and j and columns i and j) leaves both i and j non-critical. Since for increasing the antibandwidth we need to move entries away from the diagonal, candidates for swapping must be sufficiently far apart. In particular, if i is critical and there is some $k < i$ such that $i - k = ab$, then if j lies in the range

$$i - 2 * ab \leq j \leq i - 1,$$

swapping i and j will not lead to an increase in the antibandwidth. Thus j is only a swap candidate if it lies outside this range. Similarly, if i is critical and for some $k > i$ $k - i = ab$, to be a swap candidate j must lie outside the range

$$i + 1 \leq j \leq i + 2 * ab.$$

If j is a candidate for swapping with i , it is necessary to check the entries in both rows i and j to see if a swap is possible. A swap is not acceptable if one or more of the following holds:

1. $a_{ij} \neq 0$ and $|i - j| = ab$.
2. There exists l such that $a_{il} \neq 0$ and $|l - j| \leq ab$.
3. There exists k such that $a_{kj} \neq 0$ and $|k - i| \leq ab$.

If one of these holds, swapping i and j either decreases the antibandwidth or the number of critical nodes is not reduced. Each accepted swap while the antibandwidth is ab reduces the number of critical nodes by one. If the number of critical nodes becomes zero, we restart with antibandwidth $ab + 1$ and repeat the process until no further swaps can be made to reduce the number of critical nodes. The algorithm is summarized in Figure 3.1. Note that hill climbing cannot decrease the antibandwidth but may decrease the average antibandwidth.

The differences between hill climbing for reducing the bandwidth and for increasing the antibandwidth are (a) the definition of a critical node and (b) the checks that are needed for finding a suitable swap. For bandwidth reduction it is sufficient to keep track of the first and last entries in each row. For increasing

Algorithm 3: Hill climbing (HC)

```
outer: do
  Form the set  $V_c$  of critical nodes
  do until  $V_c$  is empty
    if (there are nodes  $i \in V_c$  and  $j \notin V_c$  such that
      swapping  $i$  and  $j$  leaves both non-critical) then
      swap  $i$  and  $j$  and remove  $i$  from  $V_c$ 
    else
      exit outer
    end if
  end do
end do outer
```

Figure 3.1: Hill-climbing algorithm

the antibandwidth the checking is more expensive since we must check each of the entries in rows i and j (unless the entries are in order of increasing column index but maintaining this ordering after a swap is also expensive). In our implementation, if i is critical, we swap i with the first suitable j that we find: there is no attempt to find the ‘best’ j (that is, the j that maximizes $\min_k\{|i - k| : i \neq k \text{ and } a_{jk} \neq 0\}$ and $\min_k\{|j - k| : j \neq k \text{ and } a_{ik} \neq 0\}$). This is partly because of the additional cost that locating the best j at each stage incurs but also because finding the best j does not necessarily lead to the best final antibandwidth. When looking for a swap, we search the rows in reverse order since we found this generally yielded better results.

4 Numerical experiments

We start by introducing our test problems. Our first set consists of the 24 two-dimensional meshes that are used by Duarte et al. [4]. They are constructed as the Cartesian product of two paths and optimal solutions for the antibandwidth maximization problem are known by construction (see [16]). Our second set (see Table 4.1) is taken from the University of Florida Sparse Matrix Collection [3] and comes from the DNVS, HB and AG-Monien groups. The order n ranges from 54 to 504,855. For problems with an unsymmetric sparsity structure, we work with $A + A^T$. The algorithms we propose are primarily designed for problems with underlying meshes and this has influenced our choice of test examples, although we emphasize that most of these do not have a regular rectangular grid structure (see, for example, Figure 4.1). Some of the smaller problems were chosen because they appear in the paper by Duarte et al. [4].

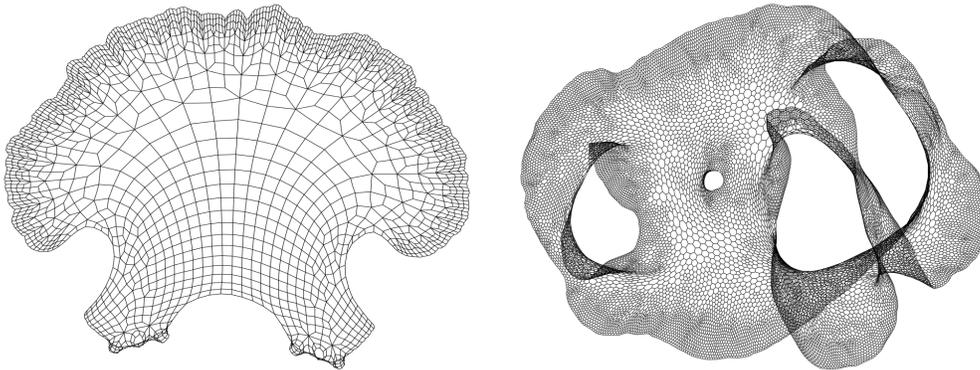


Figure 4.1: Finite-element meshes for problems AG-Monien/ukerbe1 (left) and AG-Monien/big_dual (right).

Table 4.1: Test problems. n and nz denote the order of the matrix and the number of entries in the matrix. * indicates the problem was used in [4].

Problem	n	nz	Description
HB/curtis54*	54	291	Stiff ordinary differential equation
HB/dwt_234	234	834	2D structural engineering problem
HB/saylr1	238	1 128	14×17 grid
AG-Monien/grid1	252	952	2D finite-element problem
HB/nos7	330	4 617	$9 \times 9 \times 9$ grid
HB/can.445*	445	3 809	Finite-element mesh from aircraft design
HB/nos5	468	5 172	3D finite-element approximation of building
HB/662.bus*	662	2 474	Model of power system network
HB/nos6*	675	3 255	Poisson's equation in L-shaped region
HB/saylr3	1 000	3 750	3D finite-element problem from reservoir simulation
HB/sherman4*	1 104	3 786	3D finite-element problem from oil reservoir modelling
AG-Monien/netz4504	1 961	5 156	2D finite-element problem
HB/lshp2614	2 614	17 980	Triangular finite-element mesh on 2D L-shaped region
AG-Monien/grid2	3 276	12 864	2D finite-element problem
HB/saylr4	3 564	22 316	$33 \times 6 \times 18$ grid
HB/sherman3	5 005	20 033	3D finite-element problem from oil reservoir modelling
AG-Monien/ukerbel	5 891	15 704	2D finite-element problem
AG-Monien/big_dual	30 269	89 858	2D finite-element problem
DNVS/ship_001	34 920	3 777 036	3D finite-element model of ship structure
AG-Monien/brack2	62 631	733 118	3D finite-element problem
DNVS/shipsec8	114 919	3 303 533	3D finite-element model of ship structure
DNVS/fcondp2	201 822	11 294 316	3D finite-element model of oil production platform
INPRO/msdoor2	415 863	19 173 163	Structural problem
Schenk_AFE/af_shell1	504 855	17 562 051	Structural problem

The implementations of our algorithms used in this paper are written in Fortran 95; all experiments are performed on a single core of an Intel Xeon E5620 using the gfortran compiler (version 4.4.3) with the -O3 option.

4.1 Results for the LB algorithm

To verify that the LB algorithm performs as expected for mesh problems, in Table 4.2, we present results for the 2-dimensional mesh problems used by Duarte et al. The starting node s is computed using the modified GPS algorithm of Reid and Scott [17]. The optimal solution is also given. We see that, in each case, the LB algorithm computes the optimal maximum antibandwidth or is within two of the optimal. Note that, for each mesh the reordering time was less than 10^{-3} seconds.

Table 4.2: Maximum antibandwidths for 2-dimensional mesh problems computed using the LB algorithm.

Problem	LB	Optimal	Problem	LB	Optimal
mesh9x9	36	36	mesh130x7	451	452
mesh50x2	49	49	mesh120x8	476	476
mesh34x3	49	50	mesh110x9	490	491
mesh25x4	47	48	mesh100x10	494	495
mesh20x5	47	48	mesh50x20	489	490
mesh10x10	44	45	mesh40x25	486	488
mesh17x6	47	48	mesh60x17	501	502
mesh13x8	47	48	mesh34x30	494	495
mesh15x7	49	49	mesh80x13	513	514
mesh12x9	49	49	mesh70x15	517	518
mesh11x11	55	55	mesh90x12	533	534
mesh12x12	66	66	mesh33x33	528	528

Table 4.3 reports results for the practical examples described in Table 4.1 and in the upper part of Figures 4.2 to 4.6 the sparsity patterns for some of the problems are plotted both before and after reordering using the level-based algorithm. With the exception of problems **AG-Monien/grid1** and **AG-Monien/grid2** (which have initial maximum antibandwidths of 12 and 197, respectively), the initial maximum antibandwidth is 1 but there is a large range of values for the initial average antibandwidth. We see that, for some of the problems with regular grids of square or cubic elements (such as **HB/saylr1** and **HB/nos7**), ab and av are increased substantially by relabelling and the level-based ordering gives the optimal (or close to optimal) antibandwidths (from equation (2.1), the optimal ab for **HB/saylr1** is 112 and from (2.2) for **HB/nos7** it is approximately 334). However, for other examples (including the **DNVS** problems, **HB/dwt_234** and **HB/nos5**), while the average distance between the diagonal and nearest off-diagonal entry increases (so that av increases), the maximum antibandwidth remains small. For some of these latter examples, towards the end of the relabelling, we have to give consecutive labels to nodes that are close to each other in the level-set structure; we can see this in Figure 4.5 for problem **AG-Monien/big_dual**. In the case of **HB/1shp2614**, the grid comprises triangular elements and, in general, for some r a node has neighbours in each of the level sets l_{r-1} , l_r , and l_{r+1} and so we do not have the even and then the odd level set labelling that is possible for square elements. These results show that using the level-based algorithm is not sufficient on its own.

We remark that, in the Cuthill-McKee algorithm, the unlabelled neighbours u of $v \in l_{k-1}$ are labelled in order of increasing degree. We tried modifying Algorithm 2 so that the nodes within each level set were preordered by increasing degree. We found that this did not, in general, improve the maximum antibandwidth and, for some problems, it gave poorer results.

Table 4.3: Results of applying the level-based (LB) algorithm. The maximum antibandwidth (ab) after the LB algorithm is given. The initial average antibandwidth (av_i) is also given together with $ratio_{LB} = av_{LB}/av_i$, where av_{LB} denotes the average antibandwidth after the LB algorithm.

Problem	ab	av_i	$ratio_{LB}$
HB/curtis54	4	1.15	10.5
HB/dwt_234	2	2.97	10.6
HB/saylr1	111	1.00	113.2
AG-Monien/grid1	116	82.2	1.45
HB/nos7	330	1.00	340.1
HB/can_445	5	1.19	68.0
HB/nos5	7	2.37	33.8
HB/662_bus	29	45.1	5.44
HB/nos6	329	1.00	330.4
HB/saylr3	1	317.1	1.68
HB/sherman4	257	558.5	1.23
AG-Monien/netz4504	671	18.7	51.7
HB/lshp2614	14	1.00	650.9
AG-Monien/grid2	1626	909.3	1.79
HB/saylr4	1726	1.68	1033.3
HB/sherman3	30	2111.0	1.34
AG-Monien/ukerbe1	2054	57.2	51.7
AG-Monien/big_dual	57	3.54	3400.2
DNVS/ship_001	1	1.00	564.4
AG-Monien/brack2	11	191.8	32.9
DNVS/shipsec8	3	1.00	3886.0
DNVS/fcondp2	5	1.00	8185.0
INPRO/msdoor	1	1.00	14200.0
Schenk_AFE/af_shell1	488	1.00	27360.0

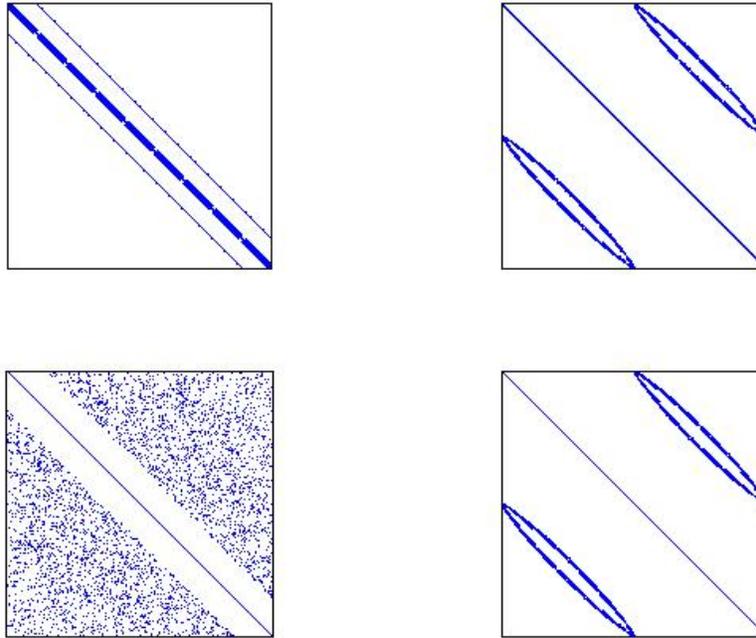


Figure 4.2: HB/nos7 (upper left), after reordering using the level-based algorithm (upper right), after hill climbing applied to the initial ordering (lower left) and to the level-based ordering (lower right).

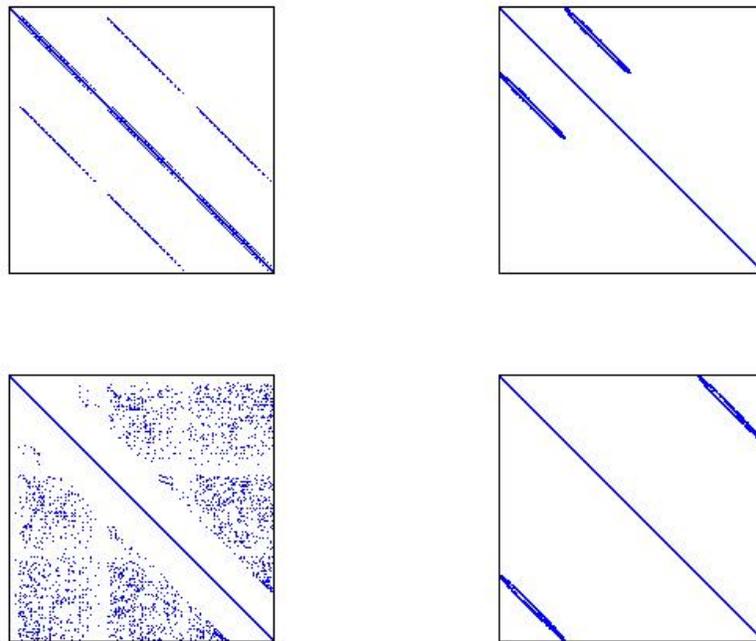


Figure 4.3: HB/sherman4 (upper left), after reordering using the level-based algorithm (upper right), after hill climbing applied to the initial ordering (lower left) and to the level-based ordering (lower right).

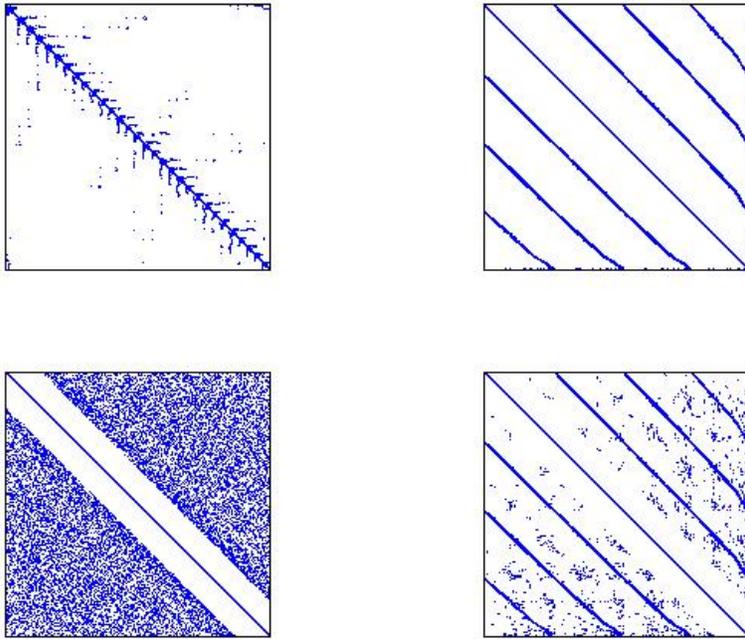


Figure 4.4: HB/1shp2614 (upper left), after reordering using the level-based algorithm (upper right), after hill climbing applied to the initial ordering (lower left) and to the level-based ordering (lower right).

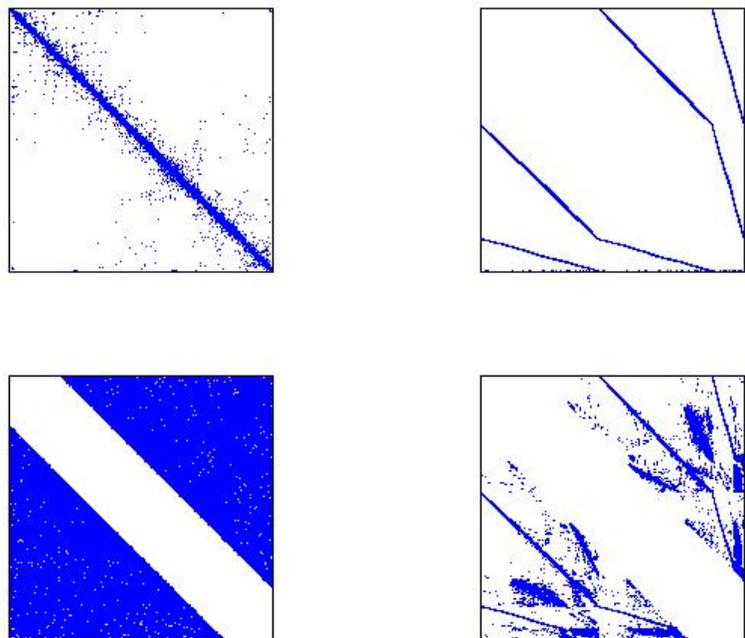


Figure 4.5: AG-Monien/big_dual (upper left), after reordering using the level-based algorithm (upper right), after hill climbing applied to the initial ordering (lower left) and to the level-based ordering (lower right).

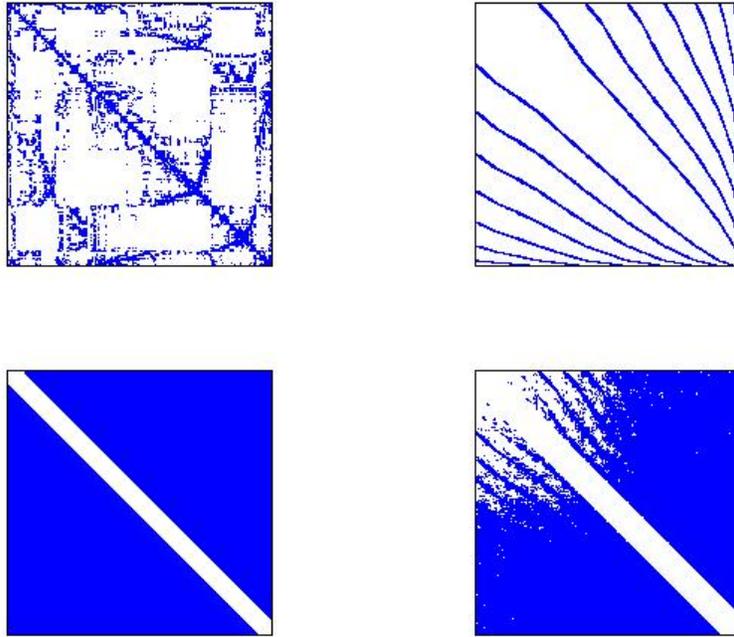


Figure 4.6: `AG-Monien/brack2` (upper left), after reordering using the level-based algorithm (upper right), after hill climbing applied to the initial ordering (lower left) and to the level-based ordering (lower right).

4.2 Hill climbing results

In Table 4.4, we report results for hill climbing applied to the initial ordering and to the level-based ordering; in the lower part of Figures 4.3 to 4.5 the sparsity patterns after hill climbing are plotted for a subset of our test problems. As already noted, for some of the grid problems (and the 2-dimensional mesh problems), the level-based ordering gives the optimal (or close to optimal) maximum antibandwidths and so hill climbing cannot improve them further. Comparing columns 2 and 4 of Table 4.4, we see that, with the exception of `HB/lshp2614`, applying HC to the level-based ordering gives a larger maximum antibandwidth than applying it to the initial ordering. This illustrates the importance of providing the hill-climbing algorithm with a good initial ordering. As expected, hill climbing applied to the level-based ordering can decrease the average antibandwidth, although the amount by which it decreases is typically less than 5 percent. We observe that the sparsity patterns that we get by applying hill climbing to the initial ordering and the level-based ordering can be very different, even when the value of ab is not too different (for example, `AG-Monien/big_dual` in Figure 4.5).

The reverse Cuthill-McKee algorithm is commonly used in place of the Cuthill-McKee algorithm since, although reversing the ordering leaves the bandwidth unchanged, it can reduce the profile [5]. Reversing the LB ordering leaves the maximum antibandwidth and the average antibandwidth unchanged. However, if hill climbing is then applied, the final antibandwidth and average antibandwidth are generally not the same. In our experiments, we found for the majority of our test cases the hill climbing results for the reverse LB ordering were poorer than for the LB ordering and so we do not recommend using the reverse ordering.

Table 4.4: The maximum antibandwidth (ab) after hill climbing is applied to the initial ordering (HC_i), after the level-based algorithm (LB) and after hill-climbing follows the level-based algorithm (LB+HC). The largest maximum antibandwidth is in bold. The ratios of the average antibandwidths to the original average antibandwidth are also given.

Problem	ab			average antibandwidth ratios		
	HC_i	LB	LB+HC	$ratio_{HC_i}$	$ratio_{LB}$	$ratio_{LB+HC}$
HB/curtis54	8	4	8	10.5	10.5	10.7
HB/dwt_234	50	2	80	28.5	10.6	33.1
HB/saylr1	45	111	111	68.4	113.2	113.2
AG-Monien/grid1	70	116	116	1.08	1.45	1.45
HB/nos7	105	330	330	159.8	340.1	340.1
HB/can_445	46	5	52	54.0	68.0	68.0
HB/nos5	32	7	49	21.5	33.8	31.8
HB/662_bus	125	29	163	4.64	5.44	5.39
HB/nos6	146	329	329	208.6	330.4	330.4
HB/saylr3	175	1	627	1.62	1.68	2.36
HB/sherman4	168	257	815	1.25	1.23	1.72
AG-Monien/netz4504	344	671	671	33.3	51.7	51.7
HB/lshp2614	343	14	337	460.8	650.9	643.0
AG-Monien/grid2	591	1626	1626	0.99	1.79	1.79
HB/saylr4	469	1726	1726	448.0	1033.3	1034.5
HB/sherman3	693	30	3509	1.27	1.34	1.96
AG-Monien/ukerbel1	1264	2054	2054	35.7	51.7	51.7
AG-Monien/big_dual	5760	57	6645	2663.0	3400.2	3451.0
DNVS/ship_001	244	1	319	347.2	564.4	534.7
AG-Monien/brack2	3480	11	4984	30.0	32.9	46.0
DNVS/shipsec8	1664	3	2246	2337.0	3886.0	3630.0
DNVS/fcondp2	2921	5	4020	4097.0	8185.0	7681.0
INPRO/msdoor	7115	1	8137	9849.0	14200.0	13910.0
Schenk_AFE/af_shell11	12970	488	14973	16980.0	27360.0	25990.0

4.3 Relaxed hill climbing

In the hill-climbing algorithm, i and j are swapped only if the swap leaves both i and j non-critical. We have investigated whether it can be advantageous to perform a swap even if it results in j becoming critical. The idea is, if hill climbing has stalled (that is, no further swaps can be made to reduce the number of critical entries), we allow a swap that leaves the number of critical entries unchanged in the hope that such a move will lead to later swaps that do reduce the critical entries. We refer to this variant as *relaxed* hill climbing. Note that when implementing this variant, care has to be taken to avoid getting into a cycle of swaps that give no gains. Furthermore, the relaxed strategy is only employed once the standard hill climbing strategy has stalled (using it from the start led to much poorer results). Results for relaxed hill climbing are given in Table 4.5 (the problems for which the relaxed strategy gave no gain are omitted). We found that the antibandwidth increased for only a few of our test problems and, in general, the additional cost of performing extra searches and swaps for relaxed hill climbing was not beneficial.

Table 4.5: A comparison of the maximum antibandwidths computed using the standard and relaxed hill-climbing strategies.

Problem	standard	relaxed
HB/can_445	52	56
HB/nos5	49	54
HB/lshp2614	337	351
AG-Monien/brack2	4984	5453

4.4 The effect of random initial permutations

Finally, we have looked at using the algorithms after applying random symmetric permutations to the given matrix. For the regular 2-dimensional mesh problems that comprise our first test set, such permutations have little effect: in all our tests, the computed maximum antibandwidth was again the optimal or within two of the optimal. The results for our more general test set are shown in Table 4.6. It is indeed the case that, if hill climbing is applied directly, better antibandwidths can often be found by considering random permutations. This is because a different initial ordering causes a different sequence of swaps to be performed. We could impose more stringent conditions on when a swap is performed to try and reduce sensitivity. The penalty would be further overheads and, in any case, the sensitivity can be viewed as advantageous as it allows us, by permuting and rerunning, to explore the “landscape” of the objective function and hopefully to then get close to finding the global optimal antibandwidth.

As we would expect, for many problems, the maximum antibandwidth obtained from the level-based algorithm is not sensitive to the initial ordering. Where there is a difference between the best and the worse LB maximum antibandwidths it is because the pseudodiameter computed by the modified GPS algorithm is dependent on the initial ordering and this in turn effects the level set structure used by the LB algorithm. Moreover, tie-breaking when selecting the nodes within each level can have an effect. Again, we could impose rules regarding tie-breaking that would reduce sensitivity.

5 Comparisons with other approaches

So far, we have presented results for our proposed level-based algorithm with hill-climbing refinement (LB+HC). In this section, we consider the GRASP and GSpectral approaches and perform some comparisons.

GRASP (Greedy Randomized Adaptive Search Procedure) is a well-known metaheuristic, that has been successfully applied to many hard combinatorial optimization problems. It is an iterative process, in which each GRASP iteration comprises two phases: construction and local search. The construction phase builds

Table 4.6: Best and worse maximum antibandwidths using the given ordering and nine random permutations; where only one number is reported, the best and worse are the same. HC_p is the antibandwidth after hill climbing is applied to the permuted matrix.

Problem	HC_p	LB	LB+HC
HB/curtis54	[5,8]	[4,6]	[7,8]
HB/dwt_234	[36,50]	2	[48,80]
HB/saylr1	[42,54]	[111,112]	[111,112]
AG-Monien/grid1	[43,70]	116	116
HB/nos7	[104,131]	330	330
HB/can_445	[40,53]	[1,5]	[47,55]
HB/nos5	[32,43]	[3,7]	[43,49]
HB/662.bus	[87,125]	[4,29]	[126,163]
HB/nos6	[114,146]	329	329
HB/saylr3	[154,193]	1	[625,627]
HB/sherman4	[158,187]	[257,258]	[815,817]
AG-Monien/netz4504	[308,438]	671	671
HB/lshp2614	[315,359]	[11,16]	[337,423]
AG-Monien/grid2	[537,662]	[1624,1626]	[1624,1626]
HB/saylr4	[441,580]	[1724,1726]	[1724,1726]
HB/sherman3	[557,755]	[29,30]	[2016,3509]
AG-Monien/ukerbe1	[971,1264]	2054	2054
AG-Monien/big_dual	[5267,6267]	[31,259]	[6526,6645]
DNVS/ship_001	[188,244]	1	319
AG-Monien/brack2	[3025,3480]	[11,44]	[4984,5313]
DNVS/shipsec8	[1552,1664]	[1,3]	[2220,2279]
DNVS/fcondp2	[2703,3035]	[1,8]	[3815,4074]

a feasible solution, whose neighborhood is explored by the local search. The best solution over all **GRASP** iterations is returned as the result. A detailed description of both the construction and local search phases for the antibandwidth maximization problem is given in Duarte et al. [4].

The **GSpectral** algorithm of Hu et al. [7] approximates the antibandwidth maximization problem by finding a permutation such that the sum of the squares of the difference of the row and column indices is maximized, that is,

$$\max_{p \in P} \sum_{a_{ij} \neq 0} (p(i) - p(j))^2. \quad (5.2)$$

Here P is the set of all possible permutations of $\{1, 2, \dots, n\}$. The vector p is a permutation and $p(i)$ is the i -th element of this vector. By relaxing this problem further so that the row and column labels are real numbers $c(i)$ with a normalization constraint, this problem becomes

$$\max_{c \in \mathbb{R}^n} \sum_{a_{ij} \neq 0} (c(i) - c(j))^2, \quad \text{subject to} \quad \sum_{k=1}^n c(k)^2 = 1. \quad (5.3)$$

It can be shown that c is the eigenvector corresponding to the largest eigenvalue of the Laplacian of the graph induced by the matrix A . Once c has been computed, a row and column permutation is obtained by ordering the entries of c . **GSpectral** then uses a greedy local refinement algorithm to try and improve the antibandwidth. This refinement proceeds by swapping the labels of pairs of indices, with a swap allowed if it increases the local antibandwidths, that is, a swap between i and j is performed provided that $\min\{ab(i), ab(j)\}$ increases, where $ab(i)$ is the local antibandwidth given by

$$ab(i) = \min_k \{|i - k| : i \neq k \text{ and } a_{ik} \neq 0\}.$$

Swapping continues until no further swaps can be found. While the eigenvector c of the Laplacian can be found in near linear time for a sparse graph using power iterations, the simple swapping procedure can take time quadratic in the dimension of the matrix, making it expensive for large problems. Note that Hu et al.'s interest was not in large problems and so the time for the refinement step was not a major concern for them. We observe also that it is possible to terminate the refinement process before it has converged (perhaps after a prescribed time limit has been reached or once the rate of progress appears to have slowed) but this was not considered by Hu et al.

In Table 5.7, we compare our computed maximum antibandwidths with those obtained using **GSpectral** and **GRASP** (with default settings); computational times are also given. We note that for some examples, the reported time for **LB+HC** is given as zero since the actual time taken is less than 10^{-3} seconds. The results show that, in some cases, our approach is successful in computing antibandwidths that are competitive with (or are larger than) those from **GSpectral** and **GRASP** but for other problems one or both of the latter gives better results. However for all problems, our algorithm is faster, significantly so on larger problems, than **GSpectral** and **GRASP**. As expected, the **LB+HC** algorithm works well on the mesh-based problems but less well on problems from other application areas (although in all cases, the **LB+HC** algorithm substantially increases the maximum antibandwidth compared to the initial ordering).

The level-based algorithm is very fast but for the large problems using hill-climbing adds a significant overhead; indeed, it accounts for almost all of the **LB+HC** time. For example, for problem **DNVS/fcondp2**, the time to run the level-based algorithm is 0.27 seconds, with the hill-climbing refinement then taking more than 102 seconds. For even larger problems, hill climbing is expensive and thus it may be necessary to consider terminating early (although we have no way of knowing whether or not the algorithm is close to converging). Similarly, for the **GSpectral** algorithm, as anticipated most of the time is in the greedy refinement phase, with the cost of the computation of the eigenvector being almost negligible by comparison. For problem **AG-Monien/big_dual**, the time to compute the eigenvector is 0.14 seconds while the greedy refinement requires 584 seconds. As expected, **GRASP** is the most expensive approach and for many of our test problems, the time limit of 1 hour that we imposed for practical reasons, was exceeded. Note that

the limit was overrun by HB/dwt_234, one of the smallest of our test examples. GRASP was left running overnight on problem AG-Monien/grid2 but still did not complete. Thus it is clear that it does not offer a method that is practical for any but small problems.

Table 5.7: A comparison of the maximum antibandwidth computed using GSpectral, GRASP and the LB+HC algorithm. The largest maximum antibandwidth is in bold. Times are also given. * indicates time limit of 1 hour exceeded.

Problem	n	nz	ab			Times (seconds)		
			GSpectral	GRASP	LB+HC	GSpectral	GRASP	LB+HC
HB/curtis54	54	291	9	12	8	0.01	0.59	0.000
HB/dwt_234	234	834	77	*	80	0.02	*	0.000
HB/saylr1	238	1 128	72	107	111	0.02	60.2	0.000
AG-Monien/grid1	252	952	77	114	116	0.02	57.7	0.000
HB/nos7	330	4 617	177	324	330	0.22	854.9	0.000
HB/can_445	445	3 809	62	84	52	0.15	547.8	0.001
HB/nos5	468	5 172	64	66	49	0.32	257.8	0.002
HB/662_bus	662	2 474	164	219	163	0.13	252.0	0.001
HB/nos6	675	3 255	212	315	329	0.16	736.9	0.000
HB/saylr3	1 000	3 750	249	*	627	0.26	*	0.010
HB/sherman4	1 104	3 786	274	*	815	0.30	*	0.006
AG-Monien/netz4504	1 961	5 156	652	*	671	1.03	*	0.002
HB/lshp2614	2 614	17 980	512	*	337	5.32	*	0.002
AG-Monien/grid2	3 276	12 864	823	*	1 626	4.69	*	0.000
HB/saylr4	3 564	22 316	873	*	1 726	6.29	*	0.000
HB/sherman3	5 005	20 033	1 246	*	3 509	8.93	*	0.37
AG-Monien/ukerbe1	5 891	15 704	1 990	*	2 054	11.92	*	0.001
AG-Monien/big_dual	30 269	89 858	10 031	*	6645	584.5	*	0.13
DNVS/ship_001	34 920	3 777 036	*	*	319	*	*	15.88
AG-Monien/brack2	62 631	733 118	*	*	4 984	*	*	6.93
DNVS/shipsec8	114 919	3 303 533	*	*	2 246	*	*	194.0
DNVS/fcondp2	201 822	11 294 316	*	*	4 020	*	*	102.5
INPRO/msdoor2	415 863	19 173 163	*	*	8 137	*	*	2 806
Schenk_AFE/af_shell1	504 855	17 562 051	*	*	14 973	*	*	2376

6 Concluding remarks and future directions

In this paper, we have discussed the antibandwidth maximization problem. We have proposed a straightforward but new approach to this practical problem that is based on extending the ideas for the bandwidth minimization problem to the antibandwidth maximization problem. The proposed LB method is inspired by the well-known Cuthill-McKee algorithm and uses the level sets of a graph. To try and improve the solution quality, we have employed a local hill-climbing algorithm. Our focus has been on problems with an underlying mesh and, for these problems, we have demonstrated that our methods are able to compute high quality orderings quickly.

A disadvantage of our current approach is that it is a simple two step approach: given an initial ordering, it computes a level-based ordering and then refines it using hill climbing. While this gives good results for many of the test problems that have an underlying mesh, it does not always work well on more general classes of problems. This suggests that we need to develop further maximization antibandwidth algorithms for non-mesh problems. In a future study, we plan to explore other techniques that are designed for bandwidth reduction to see if the ideas involved can be modified for the antibandwidth maximization problem. In particular, we will look at using a node-centroid algorithm [11, 18] combined with hill climbing and a multilevel approach (see [8] for a multilevel algorithm for reducing the profile of a symmetric matrix).

Finally, in this paper our aim is to maximize the antibandwidth and we do this using a level-based ordering. If the goal is to maximize the average antibandwidth, a different ordering may be beneficial. In the literature on profile minimization, an ordering based on the Fiedler vector has been shown to lead to better results compared to the level-based reverse Cuthill-McKee algorithm [9]. We have experimented with using the Fiedler vector to maximize the average antibandwidth; preliminary results show that for some of our test problems this can work well. We plan to investigate and other refinement algorithms this further.

Acknowledgements

We gratefully acknowledge the use of the University of Florida Sparse Matrix Collection, maintained and made available by Tim Davis and Yifan Hu. We thank Mauricio Resende for giving us access to the GRASP package. We would also like to thank the two anonymous referees whose constructive comments have helped improved the overall presentation of this paper.

References

- [1] P. CAPPANERA, *A survey of obnoxious facility location problems*, Technical Report TR-99-11, Dipartimento di Informatica, Universit a di Pisa, 1999.
- [2] E. CUTHILL AND J. MCKEE, *Reducing the bandwidth of sparse symmetric matrices*, in Proceedings of the 24th National Conference of the ACM, Brandon Systems Press, 1969.
- [3] T. DAVIS AND Y. HU, *The University of Florida Sparse Matrix Collection*, ACM Transactions on Mathematical Software, 38 (2011). Article 1, 25 pages.
- [4] A. DUARTE, R. MARTÍ, M. RESENDE, AND R. SILVA, *GRASP with path relinking heuristics for the antibandwidth problem*, Networks, (2011). doi: 10.1002/net.20418.
- [5] A. GEORGE, *Computer implementation of the finite-element method*, Technical Report STAN CS-71-208, Ph.D Thesis, Department of Computer Science, Stanford University, 1971.
- [6] N. GIBBS, W. POOLE, AND P. STOCKMEYER, *An algorithm for reducing the bandwidth and profile of a sparse matrix*, SIAM Journal on Numerical Analysis, 13 (1976), pp. 236–250.

- [7] Y. HU, S. KOBOUROV, AND S. VEERAMONI, *On maximum differential graph coloring*, in Proceedings of the 18th international conference on graph drawing (GD'10), Springer-Verlag, 2011, pp. 274–286.
- [8] Y. HU AND J. SCOTT, *A multilevel algorithm for wavefront reduction*, SIAM Journal on Scientific Computing, 23 (2001), pp. 1352–1375.
- [9] G. KUMFERT AND A. POTHEN, *Two improved algorithms for envelope and wavefront reduction*, BIT, 37:3 (1997), pp. 559–590.
- [10] J. Y.-T. LEUNG, O. VORNBERGER, AND J. WITTHOFF, *On some variants of the bandwidth minimization problem*, SIAM Journal on Computing, 13 (1984), pp. 650–667.
- [11] A. LIM, B. RODRIGUES, AND F. XIAO, *A centroid-based approach to solve the bandwidth minimization problem*, Proceedings of the 37th Hawaii international conference on system sciences, IEEE, (2004).
- [12] R. MARTÍ, M. LAGUNA, F. GLOVER, AND V. CAMPOS, *Reducing the bandwidth of a sparse matrix with tabu search*, European J. of Operational Research, 135 (2001), pp. 211–220.
- [13] Z. MILLER AND D. PRITIKIN, *On the separation number of a graph*, Networks, 19 (1989), pp. 651–666.
- [14] E. P. NANA, I. PLANA, V. CAMPOS, AND R. MARTÍ, *GRASP and path relinking for the matrix bandwidth minimization*, European J. of Operational Research, 153 (2004), pp. 200–210.
- [15] G. PAULINO, I. MENEZES, M. GATTASS, AND S. MUKHERJEE, *A new algorithm for finding a pseudoperipheral vertex or the endpoints of a pseudodiameter in a graph*, Communications in Numerical Methods in Engineering, 10 (1994), pp. 913–926.
- [16] A. RASPAUD, H. SCHRÖDER, O. SÝKORA, L. TÖRÖK, AND I. VRT'O, *Antibandwidth and cyclic antibandwidth of meshes and hypercubes*, Discrete Mathematics, 309 (2009), pp. 3541–2552.
- [17] J. REID AND J. SCOTT, *Ordering symmetric sparse matrices for small profile and wavefront*, International Journal on Numerical Methods in Engineering, 45 (1999), pp. 1737–1755.
- [18] ———, *Reducing the total bandwidth of a sparse unsymmetric matrix*, SIAM Journal on Matrix Analysis and Applications, 28 (2006), pp. 805–821.
- [19] L. TÖRÖK AND I. VRT'O, *Antibandwidth of three-dimensional meshes*, Discrete Mathematics, 310 (2010), pp. 505–510.
- [20] L. YIXEN AND Y. JINJIANG, *The dual bandwidth problem for graphs*, Journal of Zhengzhou University, 35 (2003), pp. 1–5.