STUDIA UNIV. BABEŞ–BOLYAI, INFORMATICA, Volume **LVI**, Number 3, 2011

AVERAGE BANDWIDTH REDUCTION IN SPARSE MATRICES USING HYBRID HEURISTICS

LIVIU OCTAVIAN MAFTEIU-SCAI, VIOREL NEGRU, DANIELA ZAHARIE, AND OVIDIU ARITONI

Abstract. This paper proposes a hybrid heuristic aiming the reduction of the average bandwidth of sparse matrices. The proposed heuristic combines a greedy selection of rows and columns to be interchanged with an approach based on a genetic-inspired crossover. Several ideas of limiting the search tree are also investigated. Preliminary numerical experiments illustrate the fact that the proposed heuristic leads to better results with respect to the average bandwidth than the classical Cuthill-McKee algorithm.

1. The Matrix Bandwidth Reduction Problem

Reducing the bandwidth of sparse matrices is important in solving large linear systems of equations and is especially useful in designing efficient parallel implementation of the solving procedures. The bandwidth of a matrix is related to the concentration of nonzero elements around the main diagonal and can be measured in different ways. The most frequently used bandwidth measure is the maximum of the absolute value of the difference between the row and column indices of nonzero elements, i.e. $bw = \max\{|i-j|; a_{ij} \neq 0, i = \overline{1,n}, j = \overline{1,n}\}$ for a square matrix A of size n. However, the bw measure does not provide too much information about the grouping pattern of non-zero elements around the main diagonal, which is particularly important in balancing the load on different processors in the case of parallel implementations [1]. This motivates the interest in defining and using other measures of the compactness of the non-zero elements around the main diagonal. In this work we use the average bandwidth proposed in [8] and defined by Eq. (1), where m

Received by the editors: April 10, 2011.

2010 *Mathematics Subject Classification.* 68T20, 65F30.

1998 *CR Categories and Descriptors.* I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search – *Heuristic Methods*; G.1.3 [**Numerical Analysis**]: Numerical Linear Algebra – *Sparse, structured, and very large systems (direct and iterative methods).*

Key words and phrases. bandwidth reduction, sparse matrices, heuristic algorithms.

denotes the number of non-zero elements of the matrix A.

$$mbw = \frac{1}{m} \sum_{a_{ij} \neq 0} |i - j| \tag{1}$$

If the distribution of the non-zero elements is symmetric with respect to the main diagonal, the average bandwidth can be computed by taking into account only the non-zero elements which are in the upper triangular part of the matrix, i.e. $i < j$. Minimizing the average bandwidth instead of bw has the advantage of leading to a more uniform distribution of non-zero elements around the main diagonal.

The bandwidth reduction problem consists of finding a permutation of rows and columns such that the resulting matrix has a smaller value for the bandwidth measure. Most methods designed to solve the matrix bandwidth reduction problem are based on the corresponding graph formulation: find a labeling of the vertices of a graph such that most connections are between vertices having close labels.

The methods for bandwidth matrix reduction can be grouped in two main classes: exact and heuristic ones. The exact methods are mainly based on branch and bound search for the optimal labeling of vertices (which is equivalent with an optimal permutation of rows and columns in the matrix). The branch and bound approach proposed by DelCorso and Manzini [4] is based on the concept of upper partial ordering and on an estimation of the lower bound used to prune some nodes in the search tree [1]. Caprara and Salazar [2] improved the lower bound given in [4] and introduced a control mechanism of the enumeration process by introducing new pruning criteria. A more recent approach is proposed in [9]. Its main particularity is the fact that it combines the branch and bound search with some information collected based on a heuristic search. The main disadvantage of the exact methods is related to their computational cost. This computational cost can be reduced by using heuristic techniques (which, on the other hand, do not guarantee the solution optimality). One of the best known heuristic technique, used in many comparative studies, is that proposed by Cuthill and McKee [3]. It is based on generating labelings based on some rooted level structures and has the disadvantage of being costly for some configurations and of providing solutions with a bandwidth which can be far from the optimal value. In order to improve the behavior of the Cuthill-McKee heuristic, another technique, also based on rooted level structures, has been proposed in [5]. In the last years approaches based on various metaheuristics have been proposed: tabu search [10], genetic algorithms [7], particle swarm optimization [6] etc.

---

[1] An implementation is available at www.mfn.unipmn.it/~Manzini/bandmin.

The method proposed in this paper does not use neither the graph formulation nor the rooted level structures but relies on applying successive transformations directly to the matrix. Other particularities of the proposed method are: (i) it uses the average bandwidth as primary quality measure; (ii) it combines a greedy-like selection of rows/columns to be swapped with a genetic inspired crossover; (iii) it uses an elitist selection of the active nodes in the search tree, leading to an effective pruning technique.

2. The Proposed Hybrid Heuristic

The main idea of the proposed heuristic is that of reducing the average bandwidth by applying successive perturbations based on swapping pairs of rows and corresponding pairs of columns. The key elements of the heuristic are the choice of the rows/columns to be swapped and the decision of stopping the sequence of perturbations applied to a given configuration.

Unlike most heuristic bandwidth reduction algorithms which explore a space of permutations, the proposed heuristic directly explores the space of matrices. The search tree is constructed by identifying at each step the appropriate perturbations. More specifically, for each configuration (node in the search tree) the following operations are executed:

Step 1. Find all zero elements which are placed inside the band defined by the bandwidth measure, i.e. the set $S = \{(i,j); a_{ij} = 0, i < j, |i-j| \leq imbw\}$, where *imbw* denotes the ideal average bandwidth.

Step 2. For each $(i,j) \in S$ find j' which maximizes the difference $|j'-i|$ and satisfies $|j'-i| > imbw$ and $a_{ij'} \neq 0$, construct a swap pair (i,j').

Step 3. For each swap pair selected at the previous step the corresponding rows and columns of the current matrix are swapped leading to new matrices. The bandwidth measure(s) are computed for all newly generated matrices.

Step 4. From the set of matrices constructed at the previous step is selected the subset of so-called elites consisting of all matrices having a smaller average bandwidth than the parent matrix.

Step 5. Starting from the pairs of indices used in obtaining the elite matrices, new pairs of indices are constructed by a genetic inspired crossover. More specifically if (i_1, i'_1) and (i_2, i'_2) are two pairs which led to improved matrices, then the following pairs are constructed: (i_1, i'_2) and (i_2, i'_1). The matrices corresponding to all pairs obtained by crossover are constructed and their bandwidth measures are computed.

The search tree is further developed by branching only the nodes corresponding to improved matrices (elites). The branching process can be based either on a breadth first strategy or on a depth first strategy. In both cases if all newly generated matrices from a current node have a worse structure (i.e. higher value of the average bandwidth), random swap pairs can be generated (similar to a mutation operation). This would avoid a premature stopping of the search process. On the other hand, before proceeding with the exploration of the new nodes, the elements of the elites subset generated from a given node are increasingly sorted based on the bandwidth measures (*mbw* is the primary sorting criterion and *bw* is the secondary sorting criterion). This could favor an earlier generation of good configurations.

The set of criteria used to stop the search contains: (i) finding a matrix with the desired value of the bandwidth; (ii) reaching a maximal number of steps without improvement; (iii) generating the entire search tree.

This heuristic technique for average bandwidth reduction has a structure similar to that proposed in [8] but it also has some particular characteristics: elitist selection of nodes to be further processed and generation of swap pairs using crossover and mutation operators.

3. Preliminary Results

Experiment 1. The aim of the first set of experiments was to check if the algorithm is able to reach the ideal solution (*imbw* and *ibw*), if it exists. To this end, symmetric band matrices were generated and then perturbed using a pre-established perturbation ratio. For instance, for a band matrix of size 50×50, containing 100 non-zero values, having an initial bandwidth $ibw = bw_0 = 3$ and an initial average bandwidth $imbw = mbw_0 = 1.54$, the perturbation process generated a new matrix characterized by $bw_0 = 37$ and $mbw_0 = 13.5$. Three algorithms were applied: the Cuthill-McKee algorithm, the heuristic algorithm described in [8] and the hybrid algorithm proposed in this paper. In all tested cases the proposed algorithm constructed the optimal solution and the row/columns swap operations used in the perturbation step were rediscovered by the algorithm. This experiment also illustrated the importance of sorting the nodes in the search tree. An illustration of the results obtained for a symmetric matrix of size 10×10 containing 32 nonzero elements is presented in Figure 1. As can be seen in the figure the results produced by the hybrid algorithm (Hybrid a) and Hybrid b) correspond to different branches in the search tree) are characterized by a compact grouping of non-zero elements around the main diagonal. These results also illustrate the impact on using the average bandwidth as primary selection criterion in the search process. The results obtained for several randomly generated sparse matrices having sizes between $n = 15$ and $n = 100$ are presented in Table 1. For all matrices the hybrid algorithm led to the smallest value of the average bandwidth.

Experiment 2. In this experiment we used matrices from the DWT set included in the collection Harwell-Boeing collection [2]. Some results are presented in the Table 2. The hybrid heuristic led in almost all cases to the smallest

[2]Available at http://math.nist.gov/MatrixMarket/data/Harwell-Boeing.

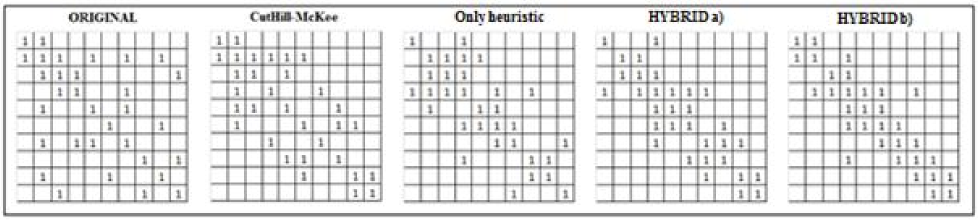

FIGURE 1. Matrices obtained by applying the Cuthill-McKee algorithm, the heuristic algorithm described in [8] and the hybrid algorithm.

TABLE 1. Results obtained for randomly generated sparse symmetric matrices (n is the matrix size, m is the number of non-zero elements.)

Matrix	Initial		Ideal case		Cuthill-McKee heuristic		Greedy heuristic [8]		Hybrid heuristic	
n/m	bw_0	mbw_0	ibw	$imbw$	bw	mbw	bw	mbw	bw	mbw
15/30	11	5.16	3	1.63	3	2.00	4	2.00	3	**1.80**
20/24	18	7.16	2	1.20	5	2.62	6	2.41	5	**2.20**
20/95	18	6.94	6	3.21	12	5.57	10	4.42	10	**4.35**
51/133	46	16.40	3	1.88	23	11.80	32	8.98	21	**6.02**
100/1650	98	34.00	19	9.32	33	38.15	78	16.13	38	**14.00**

TABLE 2. Results for sparse matrices selected from the Harwel-Boeing collection.

Matrix	Initial		Ideal case		Cuthill-McKee heuristic		Greedy heuristic [8]		Hybrid heuristic	
	bw_0	mbw_0	ibw	$imbw$	bw	mbw	bw	mbw	bw	mbw
dwt59	25	5.65	2	1.44	8	4.30	20	5.11	20	**4.09**
dwt66	44	8.34	2	1.49	3	1.88	3	1.51	4	**1.50**
dwt72	12	2.40	2	1.05	8	4.68	12	2.40	9	**2.35**
dwt87	63	19.51	3	1.86	18	**6.36**	56	10.88	31	6.40
dwt162	156	13.43	4	2.11	20	7.61	97	13.34	24	**6.65**

value of the average bandwidth. However the obtained results suggest that the effectiveness of the hybrid approach can be still improved.

4. CONCLUSIONS

The hybrid heuristic which combines the greedy choice of the pairs of rows/columns to be swapped with pairs generated by crossover starting from elite configurations proved to lead to better results than the heuristic variant proposed in [8]. The experimental results show that the proposed technique can be improved both with respect to the quality of provided solutions and with respect to its efficiency. Therefore further research will address the investigation of other strategies for exploring the search space and the parallel implementation. We will deveat the proposed technique can be improved both with respect to the quality of provided solutions and with respect to its efficiency. Therefore further research will address the investigation of other strategies for exploring the search space and the parallel implementation. We will develop in the future an recommendation system that will help us to indicate the most useful bandwidth reduction method for our matrix.

References

1. P. Arbenz, A. Cleary, J. Dongarra, and M. Hegland, *Parallel numerical linear algebra*, Nova Science Publishers, Inc., Commack, NY, USA, 2001, pp. 35–56.
2. A. Caprara and J. Salazar-González, *Laying out sparse graphs with provably minimum bandwidth*, INFORMS J. on Computing **17** (2005), 356–373.
3. E. Cuthill and J. McKee, *Reducing the bandwidth of sparse symmetric matrices*, Proc. of ACM (New York, NY, USA), ACM, 1969, pp. 157–172.
4. G. M. Del Corso and G. Manzini, *Finding exact solutions to the bandwidth minimization problem*, Computing **62** (1999), 189–203.
5. N.E. Gibbs, W.G. Poole, and P.K. Stockmeyer, *An algorithm for reducing the bandwidth and profile of a sparse matrix*, SIAM Journal on Numerical Analysis **13** (1976), 236–250.
6. A. Lim, J. Lin, and F. Xiao, *Particle swarm optimization and hill climbing for the bandwidth minimization problem*, Applied Intelligence **26** (2007), 175–182, 10.1007/s10489-006-0019-x.
7. A. Lim, B. Rodrigues, and F. Xiao, *Heuristics for matrix bandwidth reduction*, European Journal of Operational Research **174** (2006), no. 1, 69 – 91.
8. Mafteiu-Scai L.O., *Bandwidth reduction on sparse matrix*, West University of Timişoara Annals (Timişoara, Romania), vol. XLVIII fasc. 3, 2010.
9. R. Marti, V. Campos, and E. Pinana, *A branch and bound algorithm for the matrix bandwidth minimization*, European Journal of Operational Research **186** (2008), 513–528.
10. R. Marti, M. Laguna, F. Glover, and V. Campos, *Reducing the bandwidth of a sparse matrix with tabu search*, European Journal of Operational Research **135** (2001), 450–459.

West University of Timişoara, Department of Computer Science, 4 Blv. Vasile Parvan, Timişoara, România

E-mail address: {lscai,vnegru,dzaharie,oaritoni}@info.uvt.ro