Journal of Graph Algorithms and Applications
http://jgaa.info/ vol. 1, no. 1, pp. 1–1 (1111)

A Multilevel Algorithm for the Minimum 2-sum Problem

Ilya Safro *Dorit Ron* *Achi Brandt*

Faculty of Mathematics and Computer Science
The Weizmann Institute of Science
POB 26, Rehovot 76100, ISRAEL
http://www.wisdom.weizmann.ac.il
ilya.safro@weizmann.ac.il dorit.ron@weizmann.ac.il achi.brandt@weizmann.ac.il

Abstract

In this paper we introduce a direct motivation for solving the minimum 2-sum problem, for which we present a linear-time algorithm inspired by the Algebraic Multigrid approach which is based on weighted edge contraction. Our results turned out to be better than previous results, while the short running time of the algorithm enabled experiments with very large graphs. We thus introduce a new benchmark for the minimum 2-sum problem which contains 66 graphs of various characteristics. In addition, we propose the straightforward use of a part of our algorithm as a powerful local reordering method for any other (than multilevel) framework.

Article Type	Communicated by	Submitted	Revised
regular paper	Ernst W. Mayr	Oct 2004	Mar 2006

1 Introduction

The minimum 2-sum problem (M2sP) belongs to a large family of graph layout problems such as : Bandwidth, Cutwidth, Vertex Separation, Profile of a Graph, Sum Cut, etc. The M2sP appears in several applications for solving problems in the large sparse matrix computation, such as finding the minimum linear arrangement [26, 19] or the bandwidth [27]. The M2sP is also closely related to the problem of calculating the envelope size of a symmetric matrix or more precisely, to the amount of work needed in the Cholesky factorization of such a matrix [14]. In addition, the M2sP may be motivated as a model used in VLSI design, where at the placement phase it is chosen to minimize the total squared wire length [10]. Commonly for general graphs (or matrices) these problems are NP-hard and their decision versions are NP-complete [13]. The NP-completeness of the M2sP is proved in [14].

The M2sP becomes a simple quadratic optimization problem with a *known* solution, due to Hall [15], if the restriction on the solution coordinates is relaxed, i.e., the coordinates need not be all integers, as in the case where all vertices are considered to have equal unity volume (see Section 2). Hall has shown in [15] that the eigenvector v_2 which corresponds to the second smallest eigenvalue of the Laplacian of the graph (provided the graph is connected), is the best nontrivial solution to this unrestricted form of the M2sP (subject to some normalization of the solution). Arrangement of the graph vertices according to v_2 is a well known, quite successful heuristic, usually called the *spectral approach*, used for many ordering problems like the minimum linear arrangement [19], partitioning [16, 23, 22, 29], envelope reduction of sparse matrices [1], etc.

George and Pothen [14] have studied the M2sP as they used it for establishing results for the envelope reduction of matrices. They tried to evaluate the quality of the approximation for the M2sP by the spectral approach in a quantitative manner. While for some finite element graphs they indeed got close results, for general graphs the gap was profound. They suggested that this gap can be reduced by applying some local reordering (postprocessing) to the obtained results of the spectral approach.

The fact that the solution for the M2sP with real variables is extensively used as a first approximation to other graph layout problems, brings up the idea that a good solution to the *discrete* M2sP can serve as well, if not better. The first question, of course, is how well the spectral approach solves the M2sP itself, that is, how well the solution with real variables approximates the discrete setting. This question is extensively tested in our paper. In addition, since the M2sP already penalizes long distances sufficiently strongly to practically prohibit very non-uniform distribution of distances, which is typical to many other layout problems, and since its formulation is nothing but a quadratic functional, it may be considered the simplest yet central among other layout problems. As such, M2sP can and should be used (in various ways, e.g., serve as a first approximation [27]) to help solving other problems. This requires, of course, having an efficient algorithm at hand, which is exactly the purpose of our research.

In this paper we present a new multilevel algorithm for the minimum 2-sum problem based on the Algebraic MultiGrid scheme (AMG) [5, 6, 2, 9, 25, 30, 31]. The main objective of a multilevel based algorithm is to create a hierarchy of problems, each representing the original problem, but with fewer degrees of freedom. General multilevel techniques have been successfully applied to various areas of science (e.g. physics, chemistry, engineering, etc.) [4, 7]. AMG methods were originally developed for solving linear systems of equations resulting from the discretization of partial differential equations. Lately they have been applied to various other fields, yielding for example novel methods for image segmentation [28] and for the linear arrangement problem [26]. In the context of graphs it is the Laplacian matrix that represents the related set of equations. The main difference between our approach to other multilevel approaches (related to various graph optimization problems, e.g., [17]) is the coarsening scheme. While the previous approaches may be viewed as *strict* aggregation process (in which the nodes are simply blocked together into small groups and the edges are defined by the straightforward sum of the existing edges between these groups), the AMG coarsening is actually a *weighted* aggregation: each node may be divided into *fractions*, and different fractions belong to different aggregates. This enables more freedom in solving the coarser levels and avoids making hardened local decisions (such as the edge contractions made when strict aggregation is employed) before accumulating the relevant global information. The aggregation process we use here is similar to the one used for solving the minimum linear arrangement problem [26]. This part of the algorithm may, in principle, be general to many other graph layout problems (e.g., [11]), as it mainly creates the hierarchy of graphs. The diverse algorithmic ingredients, however, emerge during the disaggregation.

In the disaggregation step, the final arrangement obtained on a coarser level is projected to a finer level. This initial fine level arrangement is being further improved by applying various local reordering methods. In this article we introduce an algorithm for the strict minimization, called Window Minimization, which is based on the *simultaneous* reordering of several vertices. Then our postprocessing is intensified by Simulated Annealing (SA) [18] which is a general method to escape local minima. In the multilevel framework SA only aims at searching for *local* changes that guarantee the preservation of large-scale solution features inherited from coarser levels.

The power and robustness of our multilevel algorithm was proven by intensive experimental comparison with the spectral approach. Without the postprocessing the multilevel results are much better than the spectral ones by an average of 34.4%. After applying the same postprocessing (without simulated annealing) to both the multilevel and the spectral first approximations, the gap between the two frameworks was significantly reduced, but still the multilevel results are better by an average of 4.7%. Different parts of the postprocessing were enabled step-by-step in order to show the gradual improvement of the results. However, not only the results of the multilevel algorithm are better, but while our algorithm performs in linear time, the spectral approach is sensitive to the obtained accuracy; the trade-off between the complexity and the high ac-

curacy of the calculation of v_2 is discussed in Section 4. Our experiments show that the Algebraic Multilevel approach can be used as a first approximation for the M2sP to obtain high quality results in linear time, while the postprocessing can actually improve these results and can serve as a tool for improving any first approximation of the M2sP obtained by other methods. The implemented algorithm can be obtained at *http://www.wisdom.weizmann.ac.il/∼safro/min2sum*.

The problem definition and its generalization are described in Section 2. The multilevel algorithm along with additional optimization techniques are presented in Section 3. A comparison of our results with the spectral approach is finally summarized in Section 4.

2 Problem definition and generalization

Given a weighted graph $G = (V,\ E)$, where $V = \{1, 2, ..., n\}$, denote by w_{ij} the non-negative weight of the edge ij between nodes i and j (if $ij \notin E$ then $w_{ij} = 0$). The purpose of the minimum 2-sum problem is to find a permutation π of the graph nodes such that the cost $\sigma_2(G, \pi) = \sum_{ij} \left(w_{ij}(\pi(i) - \pi(j))^2\right)$ is minimal. In the generalized form of the problem that emerges during the multilevel solver, each vertex i is assigned a *volume* (or *length*), denoted v_i. The task now is to minimize the cost $\sigma_2(G, x, \pi) = \sum_{ij} \left(w_{ij}(x_i - x_j)^2\right)$, where $x_i = \frac{v_i}{2} + \sum_{k,\pi(k)<\pi(i)} v_k$, i.e., each vertex is positioned at its center of mass capturing a segment on the real axis which equals its length. Note that the difference $x_i - x_j$ contains the volumes of the vertices between i and j and half the volumes of i and j. The original form of the problem and the general form with equal-volume vertices are both minimized by the same permutation.

We are not interested in the worst possible cases, which are often very artificial. Our focus is on practical high-performance algorithm that will yield (in most practical cases) a good approximation to the optimum at low computational cost. Typically, the multilevel algorithms exhibit linear complexity, i.e., the computational cost in most practical cases is proportional to $|V| + |E|$.

3 The algorithm

In the multilevel framework a hierarchy of decreasing size graphs : $G_0, G_1, ..., G_k$ is constructed. Starting from the given graph, $G_0 = G$, create by *coarsening* the sequence $G_1, ..., G_k$, then solve the coarsest level directly, and finally uncoarsen the solution back to G. This entire process is called a *V-cycle*.

As in the general AMG setting, the choice of the coarse variables (aggregates), the derivation of the coarse problem which approximates the fine one and the design of the coarse-to-fine disaggregation (uncoarsening) process are all determined automatically as described below.

3.1 Coarsening: Weighted Aggregation

The coarsening used here is similar to the process we have used in solving the minimum linear arrangement problem [26]. However, for the completeness of this article, we briefly repeat it.

The coarsening is interpreted as a process of *weighted aggregation* of the graph nodes to define the nodes of the next coarser graph. In *weighted* aggregation each node can be divided into *fractions*, and different fractions belong to different aggregates. The construction of a coarse graph from a given one is divided into three stages: first a subset of the fine nodes is chosen to serve as the *seeds* of the aggregates (the nodes of the coarse graph), then the rules for interpolation are determined, thereby establishing the fraction of each non-seed node belonging to each aggregate, and finally the strength of the connections (or edges) between the coarse nodes is calculated.

Coarse Nodes. The construction of the set of seeds C and its complement, denoted by F, is guided by the principle that each F-node should be "strongly coupled" to C. Also, we will include in C nodes with exceptionally large volume, or nodes expected (if used as seeds) to aggregate around them exceptionally large volumes of F-nodes. To achieve these objectives, we start with an empty set C, hence $F = V$, and then sequentially transfer nodes from F to C, employing the following steps. As a measure of how large an aggregate seeded by $i \in F$ might grow, define its *future-volume* ϑ_i by

$$\vartheta_i = v_i + \sum_{j \in V} v_j \cdot \frac{w_{ji}}{\sum\limits_{k \in V} w_{jk}} \quad . \tag{1}$$

Nodes with future-volume larger than η times the average of the ϑ_i's are first transferred to C as most "representative". (In our tests $\eta = 2$). The insertion of additional fine nodes to C depends on a threshold Q (in our tests $Q = 0.4$) as specified by Algorithm 1. That is, a fine node i is added to C if its relative connection to C is not strong enough, i.e., smaller than Q. Also, vertices with larger values of ϑ_i are given higher priority to be chosen to belong to C.

Algorithm 1: CoarseNodes($Parameters : Q,\ \eta$)

$C \leftarrow \emptyset,\ F \leftarrow V$
Calculate ϑ_i for each $i \in F$, and their average $\overline{\vartheta}$
$C \leftarrow$ nodes i with $\vartheta_i > \eta \cdot \overline{\vartheta}$
$F \leftarrow V \setminus C$
Sort F in descending order of ϑ_i
Go through all $i \in F$ in descending order of ϑ_i
  If $\left(\sum\limits_{j \in C} w_{ij} / \sum\limits_{j \in V} w_{ij} \right) \leq Q$ **then** move i from F to C
Return C

For convenience we are currently using a library $O(n \log(n))$ sorting algorithm. However, since no exact ordering is really needed, this can be replaced by a

rough bucketing sort which has $O(n)$ complexity. We have actually implemented a simple bucketing sort and compared its results with the exact sort ones. Since no significant differences were observed, we may state that the exact sorting does not play a role in the algorithm. The only important task is to identify the vertices with exceptionally large future-volume, which can easily be achieved by $O(n)$ procedure.

The Coarse Problem. Each node in the chosen set C becomes the seed of an aggregate that will constitute one coarse level node. Define for each $i \in F$ a coarse neighborhood $N_i = \{j \in C,\ w_{ij} \geq \alpha_i\}$, where α_i is determined by the demand that $|N_i|$ does not exceed the allowed coarse neighborhood size r chosen to control complexity. (For typical values of r consider the Appendix). The classical AMG interpolation matrix P (of size $|V| \times |C|$) is defined by

$$P_{ij} = \begin{cases} w_{ij} / \sum\limits_{k \in N_i} w_{ik} & \text{for } i \in F,\ j \in N_i \\ 1 & \text{for } i \in C,\ j = i \\ 0 & \text{otherwise} \quad . \end{cases} \tag{2}$$

P_{ij} thus represents the likelihood of i to belong to the j-th aggregate. Let $I(k)$ be the ordinal number in the coarse graph of the node that represents the aggregate around a seed whose ordinal number at the fine level is k. Following the weighted aggregation scheme used in [28], the edge connecting two coarse aggregates, $p = I(i)$ and $q = I(j)$, is assigned with the weight $w_{pq}^{(coarse)} = \sum_{k \neq l} P_{ki} w_{kl} P_{lj}$. The volume of the i-th coarse aggregate is $\sum_j v_j P_{ji}$. Note that during the process of coarsening the total volume of all vertices is conserved.

Solving the coarsest level, which consists of no more than 8 nodes (otherwise a still coarser level would be introduced for efficiency) is performed directly by simply trying all possible arrangements.

3.2 Disaggregation (uncoarsening)

Having solved a coarse problem, the solution to the next-finer-level problem is initialized by first placing the seeds according to the coarse order and then adjusting all other F-nodes while aiming at the minimization of the quadratic arrangement cost. This approximation is subsequently improved by several *relaxation* (local reordering) sweeps, first compatible, then regular with or without additional stochastic elements, as explained below and summarized in Algorithm 3.

3.2.1 Initialization

Given is the arrangement of the coarse level aggregates in its generalized form, where the center of mass of each aggregate $j \in C$ is positioned at $x_{I(j)}$ along the real axis. We begin the initialization of the fine level arrangement by letting each seed $j \in C$ inherit the position of its respective aggregate: $y_j = x_{I(j)}$. At each stage of the initialization procedure, define $V' \subset V$ to be the subset of nodes that have already been placed, so we start with $V' = C$. Then proceed

by positioning each fine node $i \in V \setminus V'$ at the coordinate y_i in which the cost of the arrangement, at that moment when i is being placed, is minimized. The sequence in which the nodes are placed is roughly in decreasing order of their *relative* connection to V', that is, the nodes which have strong connections to V' compared with their connections to V are placed first. To be precise, the coordinate y_i is located at its minimum (volumes are not taken into account)

$$y_i = \frac{\sum_{j \in V'} y_j w_{ij}}{\sum_{j \in V'} w_{ij}}. \tag{3}$$

Then $V' \leftarrow V' \cup \{i\}$ and the process continues until $V' = V$. Finally each position y_i is changed to

$$x_i = \frac{v_i}{2} + \sum_{y_k < y_i} v_k \quad , \tag{4}$$

thus retaining order while taking volume (length) into account.

3.2.2 Relaxation

The arrangement obtained after the initialization is a first feasible solution for M2sP which is then improved by employing several sweeps of *relaxation*, first *compatible* then *Gauss-Seidel*-like (GS). These two types of relaxation are very similar to the above initialization: The compatible relaxation, motivated in [3], improves the positions of (only) the F-nodes according to the minimization criterion (3) (where $V' = V$) while keeping the positions of the seeds (C-nodes) unchanged. The GS relaxation is similarly performed, but for *all* nodes (including C). Each such sweep is again followed by (4).

3.2.3 Window Minimization

The cost of the arrangement can be further reduced by *strict minimization*, a sequence of rearrangements that accepts only changes which decrease the arrangement cost. Since done in the multilevel framework, it can be restricted at each level to just *local* changes, i.e., reordering small sets of neighboring nodes, which are adjacent (or relatively close) to each other at the current arrangement. It is easy to see that switching positions between several adjacent nodes is indeed a local operation, since the resulting new arrangement cost can be calculated only at the vicinity of the adjustment and not elsewhere. Such a node by node minimization was applied in our algorithm for the Minimum Linear Arrangement problem (see [26]). This method may also be used for M2sP. However, we would like to propose a more advanced method of local minimization, called *Window Minimization* (WM), which is suitable for both the multilevel and the spectral approach frameworks. The difference between WM and simple node by node minimization is that WM *simultaneously* minimizes the arrangement cost of several nodes. Given a current approximation $\tilde{x}$ to the arrangement of the graph, denote by δ_i a *correction* to $\tilde{x}_i$. Let $\mathfrak{W} = \{i_1 =$

$\pi^{-1}(p+1), ..., i_q = \pi^{-1}(p+q)\}$ be a *window* of q sequential vertices in the current arrangement, i.e., the nodes positioned at q subsequent coordinates $\tilde{x}_{i_1}, ..., \tilde{x}_{i_q}$. The local energy minimization problem associated with a given window $\mathfrak{W}$ can be formulated as follows :

$$\text{minimize } \sigma_2(\mathfrak{W}, \tilde{x}, \pi, \delta) = \sum_{i,j \in \mathfrak{W}} w_{ij}(\tilde{x_i}+\delta_i-\tilde{x_j}-\delta_j)^2 + \sum_{i \in \mathfrak{W},\ j \notin \mathfrak{W}} w_{ij}(\tilde{x_i}+\delta_i-\tilde{x_j})^2. \tag{5}$$

To prevent the possible convergence of many coordinates to one point, and, more precisely, to express the aim of having $\{x_i+\delta_i\}_{i \in \mathfrak{W}}$ an approximate permutation of $\{x_i\}_{i \in \mathfrak{W}}$, we have added the following constraints

$$\sum_{i \in \mathfrak{W}} (\tilde{x_i} + \delta_i)^m v_i \ = \ \sum_{i \in \mathfrak{W}} \tilde{x_i}^m v_i \ , \quad m = 1, 2$$

where for $m = 2$ we have neglected the quadratic term in δ_i. Note that the sums $\sum_{i \in \mathfrak{W}} \tilde{x_i}^m v_i$ for $m = 1, 2$ are invariant under permutations. Using Lagrange multipliers, the final formulation of the window minimization problem is :

$$\text{minimize } \sigma_2(\mathfrak{W}, \tilde{x}, \pi, \delta, \lambda_1, \lambda_2) = \sigma_2(\mathfrak{W}, \tilde{x}, \pi, \delta) + \lambda_1 \sum_{i \in \mathfrak{W}} \delta_i v_i + \lambda_2 \sum_{i \in \mathfrak{W}} \delta_i v_i \tilde{x_i} \quad , \tag{6}$$

under the second and third constraints of (7) below, yielding the following system of equations:

$$\begin{cases} \sum\limits_{j \in \mathfrak{W}} w_{ij}(\delta_i - \delta_j) + \delta_i \sum\limits_{j \notin \mathfrak{W}} w_{ij} + \lambda_1 v_i + \lambda_2 v_i \tilde{x_i} = \sum\limits_{j} w_{ij}(\tilde{x}_j - \tilde{x}_i) \quad i \in \{1 \dots q\} \\ \sum\limits_{i} \delta_i v_i = 0 \\ \sum\limits_{i} \delta_i v_i \tilde{x}_i = 0 \qquad . \end{cases} \tag{7}$$

Usually in a correct multilevel framework, the changes δ_i are supposed to be relatively small since the global approximation for the arrangement is inherited from the coarser levels. Their smallness is effected by the very restriction of the minimization to one window at a time. Because of the continuous formulation (7) of the problem within a window, the solution will almost always tend to move the vertices away from their initial ordering (adapted from the discrete arrangement). These changes may introduce some overlap between the vertices, but at the same time decrease the energy cost (6). It is thus expected that the δ_i's will not vanish even at the global minimum. After solving the system (7), every vertex $i \in \mathfrak{W}$ is thus positioned at $y_i = \tilde{x_i} + \delta_i$. Feasibility with respect to the volumes of the nodes is retained by applying (4). Since the size and location of $\mathfrak{W}$ are quite arbitrary, the energy cost of the new sub-arrangement is further improved by GS relaxation sweeps applied to an *enlarged* $\mathfrak{W}$, where, say 5% of the window's size at each end (if possible) are added to $\mathfrak{W}$. As usual, each sweep is followed by (4). The final obtained energy cost is then compared with the one prior to all the window changes, the minimum of the two is accepted,

updating $\tilde{x}$. We have observed actual decrease in the energy cost in about 5% of the windows.

A sweep of WM with a given window size q consists of a sequence of overlapping windows, starting from the first node in the current arrangement and stepping by $\lfloor \frac{q}{2} \rfloor$ for each additional window. One such sweep is employed for every given q, while a small number of different q's is used (in our tests there never was a need for more than 6). Our experiments show that the algorithm is robust to changes in the chosen q's; for complete details consider $WinSizes$ in the Appendix. Note, however, that q should be small enough to still guarantee linear execution time of the entire algorithm. The WM is summarized in Algorithm 2.

Algorithm 2: WindowMinimization(graph G, current order $\tilde{x}$)

Parameters: $WinSizes$, k_2 (for chosen values, consider the Appendix)

For each $q \in WinSizes$
  For $i = 1$ **To** $|V|$ **Step** $i = i + \lfloor \frac{q}{2} \rfloor$
    $\mathfrak{W} = \{\pi^{-1}(i), ..., \pi^{-1}(i + q - 1)\}$
    Solve the system of equations (7)
    Apply k_2 sweeps of GS relaxation on the enlarged $\mathfrak{W}$ with $\tilde{x} + \delta$
    $\tilde{x} \leftarrow \tilde{x} + \delta$ if the cost of the arrangement was decreased
Return $\tilde{x}$

3.2.4 Simulated Annealing

A general method to escape false local minima and advance to lower costs is to replace the strict minimization by a process that still accepts each candidate change which lowers the cost, but also assigns a positive probability for accepting a candidate step which increases the cost of the arrangement. The probability assigned to a candidate step is equal to $exp(-\Delta/T)$, where $\Delta > 0$ measures the *increase* in the arrangement cost and $T > 0$ is a temperature-like control parameter which is gradually decreased toward zero. This process, known as *Simulated Annealing* (SA) [18], in large problems would usually need to apply *very gradual* cooling (decrease of temperatures), making it extremely slow and inefficient for approaching the global optimum.

In the multilevel framework, however, the role of SA is somewhat different. At each level it is assumed that the *global* arrangement of aggregates has been inherited from the coarser levels, and thus only *local*, small-scale changes are needed. For that purpose, we have started at relatively high T, lowered it *substantially* at each subsequent sweep, until window minimization is employed.

In particular, $2k + 1$ candidate locations are examined for each vertex, each corresponds to moving it some distance l, $0 < |l| \leq k$. The initial temperature $T = T(|l|) > 0$ is calculated apriori for each distance l by aiming at the acceptance of a certain percent of changes (for instance 60%). In detail,

the probability of moving a vertex l positions ($l = \pm 1, ..., \pm k$) is $Pr(l) = z \cdot min(1, exp(-\Delta(l)/T(|l|))$, where z is a normalization factor calculated by the demands $\sum_{l=-k}^{k} Pr(l) = 1$ and $Pr(0) = z \cdot min_{l=\pm 1,...,\pm k}(1 - Pr(l)/z)$. In each additional sweep $T(|l|)$ is reduced by a factor, such as 0.6. Typically only three such cooling steps are used.

Repeated heating and cooling is successively employed for better search over the local landscape. This search is further enhanced by the introduction of a "memory"-like tool consisting of an additional permutation which stores the *Best-So-Far* (BSF) observed arrangement, which is being occasionally updated by a procedure called *Lowest Common Configuration* (LCC) [8]. LCC enables the systematic accumulation of *sub*-permutations over a sequence of different arrangements, such that each BSF sub-permutation exhibits the best (minimal) sub-order encountered so far. The cost of the obtained BSF is at most the lowest cost of all the arrangements it has observed, and usually it is lower. The use of LCC becomes more important for larger graphs, where it is expected that the optimum of a subgraph is only weakly dependent on other subgraphs. Due to the LCC procedure, it is not necessary to wait in the stochastic annealing process until all minimal sub-permutations are *simultaneously* obtained, which may take exponential time; instead it is sufficient to obtain each such minimal sub-order just once, since henceforth it is guaranteed to appear in the BSF. In detail, the BSF (of a certain level) is initialized by the arrangement obtained at the end of the strict minimization. Then the BSF is improved by the LCC procedure which updates parts of it taken from the new arrangements reached right after each heating-cooling cycle. All these accumulated updates are thus stored at the BSF, which thus represents the current calculated minimum. The complete description of the LCC algorithm is given in [26].

The entire disaggregation procedure is summarized below in Algorithm 3. The Algorithm is divided into two parts: the first approximation and the post-processing corresponding to the results supported later.

Algorithm 3: Disaggregation(coarse level $\mathcal{C}$, fine level $\mathcal{F}$)

Parameters: $k_1, ..., k_5, \gamma$ (for chosen values consider the Appendix.)

FIRST APPROXIMATION :
  Initialize $\mathcal{F}$ from $\mathcal{C}$
  Apply k_1 sweeps of compatible relaxation on $\mathcal{F}$
POSTPROCESSING :
  Apply k_2 sweeps of GS relaxation on $\mathcal{F}$
  Apply Window Minimization on $\mathcal{F}$
  Initialize $BSF \leftarrow$ current arrangement of $\mathcal{F}$
  Do k_3 cycles of heating and cooling
    Calculate $T(|l|)$ for $l = 1, ..., k_4$
    Do k_5 steps
      Apply SA within distance k_4 on $\mathcal{F}$
      Decrease all $T(|l|)$ by a factor γ

Apply Window Minimization on $\mathcal{F}$
$BSF = LCC(BSF,$ current arrangement of $\mathcal{F})$
Return BSF

4 Results and Related Works

We have implemented and tested the algorithm using standard C++, LAPACK++ [24] and LEDA libraries [20] on Linux 2.4GHz machine. The implementation is non-parallel and has not been optimized. The results (order costs and running times) should only be considered qualitatively and can certainly be further improved by more advanced implementation.

We have found only one article [14] with an implemented algorithm and numerical results for M2sP. The algorithm is based on the spectral approach. Since this test suite is relatively small to provide enough information regarding M2sP, we have launched a new, much larger test suite which consists of 66 graphs from different areas [12, 21], see Table 1. These graphs are divided into two groups according to their size : the results for the smaller ones are introduced in Tables 2 and 3, while those for the larger ones in Table 4. For all the graphs in Tables 2 and 3 we compare our results with those of the spectral approach. The numbers in columns 4-5 (marked by "**ML**" and "**ML+GS**") and 7-11 are in percentage above the cost energy presented at the column "**Quick**" (e.g., the 0.8 appearing for the first graph gd96c in column "**ML**" means that the initial cost energy is $3455 \cdot 1.008$). The first approximation obtained by the multilevel V-cycle, i.e., the arrangement obtained right after applying the compatible relaxation at the finest level is introduced in the column "**ML**" of Tables 2 and 3. We run the algorithm 100 times (using the parameters specified in the Appendix for the "**Quick**" V-cycle), each starts from a different permutation of the nodes. The best obtained results are presented here. The means of the 100 runs are worse than the corresponding "**Quick**"-values by an average of 0.51%, while the standard deviation (around the means) is 0.69% on the average. The "**ML**" results should be compared to the spectral approach results at column "**SP**" obtained by calculating the second eigenvector of the Laplacian[1] of the graph using MATLAB routine. It shows that the "**ML**" algorithm provides much better results, shorter by an average of +34.4% (excluding the statistics of bintree10, in which the improvement is much larger). Only in one case, the 10-dimensional hypercube, the spectral approach provided a lower cost of -2.8%. However, not only the obtained results are much worse. Even if the spectral method leads to the correct order, the calculations must be performed with very

[1]The algebraic representation of a graph is given by its *Laplacian* A (a $|V| \times |V|$ matrix), whose terms are defined by

$$A_{ij} = \begin{cases} -w_{ij} & \text{for } ij \in E,\ i \neq j \\ 0 & \text{for } ij \notin E,\ i \neq j \\ \sum_{k \neq i} w_{ik} & \text{for } i = j \end{cases} . \tag{8}$$

high accuracy. In fact, the precision of the second eigenvector coordinates must be at least $O(\log |V|)$ and usually much better. This is not a trivial task while one uses some approximation algorithm. Our results for the spectral approach were thus obtained with 16-digits precision of an exact algorithm. The experiments with lower precision or with approximation algorithms gave much poorer results. For example, for the three graphs bcspwr10, airfoil1 and bcsstk38 the spectral costs with 5-digits precision were 4.40E+07, 4.49E+07 and 1.20E+10 respectively, while increasing the precision to 7-digits gave 1.64E+07, 1.93E+07 and 4.40E+09. The complexity of an exact calculation of the second smallest eigenvector is $O(|V|^3)$ while the multilevel algorithm is linear in the number of edges.

We have next tested the outcome of our postprocessing on both initial sets of results. Most significant improvement was introduced by applying the Gauss-Seidel-like relaxation, as can be seen in Tables 2 and 3 column "**ML+GS**" for the multilevel algorithm and "**+GS**" for the spectral approach. The gap between the two has been reduced, but the spectral approach still provides worse results on the average by 7.1%. Next we have applied the window minimization which concludes our, so called "**Quick**" V-cycle. Comparing columns "**Quick**" with the corresponding "**+WM**" shows that the multilevel results remain better, the spectral ones are worse by an average of 4.7%.

Finally, we introduced randomness by applying Simulated Annealing. In the multilevel framework, the SA enters at all levels of the V-cycle. We refer to this version as the "**Extended**" V-cycle (its complete parameters are given in the Appendix). While the "**Quick**" V-cycle is aimed at achieving fast performance, the "**Extended**" V-cycle runs longer but succeeds in finding lower cost arrangements on the average by 1%. The means of the 100 runs of the "**Extended**" V-cycles are worse than the corresponding "**Quick**"-values by an average of 0.49% and the average of the calculated standard deviations (around the means for 100 runs) of the "**Extended**" V-cycle is 0.66%. We may conclude that the "**Extended**" V-cycle is not really needed. Almost identical results are already obtained by the "**Quick**" V-cycle, where the improvement is neither significant nor consistent, i.e., it is just within the typical standard deviation. In column "**+SA**" of Tables 2 and 3 we present the results obtained after adding SA to the spectral approach followed by the above postprocessing. The improvement is again of only 1%. Our last test was to run a very long SA after the postprocessing with the spectral approach, aiming at achieving comparable amount of work to 100 "**Extended**" V-cycles. These results are given in column "**HSA**". While improvement is naturally observed, the results on the average remain worse by about 2%, while for 6 graphs out of 37 it is worse by more than 5%.

In addition, we present the spectral lower bounds [14] for the smaller graphs (see Tables 2 and 3, column "**LB**") and calculate the gap (see Tables 2 and 3, column "$\mathbf{\Delta_{Quick}}$") between our results and the spectral lower bound. In spite of the fact that it is impossible to judge which costs are closer to the real minima, we may state that no significant indications of the existence of these lower costs were observed: in 16 out of 37 graphs our results were within 25% of the lower bounds, but on the average they were 75.4% longer.

To enrich our test suite, we present in Table 4 our "**Quick**" V-cycle results for additional 29 relatively large graphs. No spectral approach results are provided since we were not able to run (on the computers available to us) the MATLAB routine and calculate the needed eigenvector. Each result is again the best observed out of 100 runs, for which the means for 100 runs are worse than the corresponding "**Quick**"-values by an average of 0.55% and the average standard deviation is 0.47%.

The running time of the algorithm is presented in Table 6. In column "$\mathbf{T_{Quick}}$" we present the running time of one V-cycle which corresponds to the "**Quick**" column in Tables 2, 3 and 4. The running time of the suggested postprocessing added after the spectral ordering was measured for the graphs from Tables 2, 3 and we show it in column "$\mathbf{T_{Post}}$". The running time of the postprocessing corresponds to the values introduced in the "**+SA**" column of Tables 2 and 3. The dash notation ("-") corresponds to the graphs from Table 4 that were too large for usual MATLAB spectral calculation routines. In both cases the running time is measured in minutes. The star notation ("*") can be interpreted as "less than one second".

5 Conclusions

We have presented a multilevel algorithm for the minimum 2-sum problem for general graphs. The algorithm is based on the general principle that during coarsening each vertex may be associated to more than just one aggregate according to some "likelihood" measure. The uncoarsening initialization, which produces the first arrangement of the fine graph nodes, strongly relies on energy considerations (unlike usual interpolation in classical AMG). This initial order is further improved by Gauss-Seidel-like relaxation, window minimization and possibly by employing randomness, i.e., simulated annealing. The running time of the algorithm is linear, thus it can be applied to very large graphs.

We have compared our results to those obtained by the spectral approach. The calculation of the second eigenvector of the Laplacian of the graph has to be of high accuracy to provide reasonable results. Such a direct computation is of complexity $O(|V|^3)$. Still, the obtained results are much worse than the initial results obtained by our multilevel V-cycle by 34.4% on the average for the smaller sized test suite. In addition, we have applied postprocessing to both initial arrangements. The Gauss-Seidel-like relaxation improves both results most significantly. The window minimization further reduces the arrangement cost for some graphs. The final results show that the multilevel framework achieves better results of 4.7% on the average. Finally, we have added stochastisity to both algorithms. Both results were improved by about 1%. We have also tried to apply a very long SA to the final results of the postprocessing of the spectral approach. Many results have been further improved, however, some graphs (6 out of 37) still present results higher by more than 5%.

Our main conclusion is that the average and the best results of our V-cycles are better than the results of the spectral approach. We recommend

our multilevel algorithm as a general practical method for solving the minimum 2-sum problem and as a fast and high-quality method for obtaining first approximation for it. The implemented algorithm can be obtained at *http://www.wisdom.weizmann.ac.il/∼safro/min2sum*.

Appendix: Parameters

In order to control the running time of the algorithm it is important to decrease the total number of edges of the constructed coarse graphs. This is achieved by the following two parameters: the maximum allowed coarse neighborhood size r, which restricts the allowed size $|N_i|$ of the coarse neighborhood of a vertex $i \in F$ by deleting the weakest w_{ij}, $j \in C$; and the edge filtering ϵ threshold, which deletes every *relatively* weak edge ij (from the created coarse graph) if both $w_{ij} < \epsilon \cdot s_i$ and $w_{ij} < \epsilon \cdot s_j$, where $s_i = \sum_k w_{ik}$.

These two parameters and five others which control the uncoarsening procedure (see Algorithm 3) are presented in Table 5 for the "**Quick**" and "**Extended**" V-cycles we have used. The last two parameters within the SA (of Algorithm 3) were constantly chosen to be $k_5 = 4$ and $\gamma = 0.6$.

It is however important to mention that these parameters are the ones used only for the finest levels. As the coarse graphs become much smaller they are accordingly increased. This hardly affects the entire running time of the algorithm but systematically improves the obtained results. In the last column of Table 5 we specifically describe the increase introduced for each parameter as a function of level L, which usually depends on the ratio $R = max(1, |E_0|/|E_L|)$ measuring the relative decrease of the number of edges at level L compared with the original graph.

We tested many options for the window sizes in Algorithm 2. Usually these sizes were relatively small and very robust to changes. In our implementation we used $WinSizes = \{5, 10, 15, 20, 25, 30\}$, however similar results were obtained with other sets of windows, for example, $WinSizes = \{5, 9, 17, 23, 29\}$.

Acknowledgments

This research was supported by a Grant from the German-Israeli Foundation for Scientific Research and Development (G.I.F.), Research Grant Agreement No. I-718-135.6/2001, and by the Carl F. Gauss Minerva Center for Scientific Computation at the Weizmann Institute of Science.

Table 1: Benchmark for the minimum 2-sum problem.

Graph	**|V|**	**|E|**	**Graph**	**|V|**	**|E|**
gd96c	65	125	**nasa1824**	1824	18692
gd95c	62	144	**randomA2**	1000	24738
gd96b	111	193	**nasa2146**	2146	35052
gd96d	180	228	**bcsstk13**	2003	40940
dwt245	245	608	**whitaker3**	9800	28989
bintree10	1023	1022	**zcrack**	10240	30380
bus685	685	1282	**shuttleeddy**	10429	46585
bus1138	1138	1458	**randomA3**	1000	49820
gd96a	1096	1676	**nasa4704**	4704	50026
can445	445	1682	**bcsstk24**	3562	78174
c1y	828	1749	**bcsstk38**	8032	173714
c2y	980	2102	**finan512**	74752	261120
bcspwr08	1624	2213	**bcsstk33**	8738	291583
bcspwr09	1723	2394	**bcsstk29**	13830	302424
c5y	1202	2557	**ocean**	143437	409593
jagmesh1	936	2664	**tooth**	78136	452591
c3y	1327	2844	**mrng1**	257000	505048
c4y	1366	2915	**bcsstk37**	25503	557737
dwt918	918	3233	**msc23052**	23052	559817
dwt1007	1007	3784	**bcsstk36**	23052	560044
jagmesh9	1349	3876	**bcsstk31**	35586	572913
can838	838	4586	**msc10848**	10848	609464
randomA1	1000	4974	**ferotor**	99617	662431
hc10	1024	5120	**bcsstk35**	30237	709963
can1054	1054	5571	**598a**	110971	741934
can1072	1072	5686	**bcsstk32**	44609	985046
randomG4	1000	8173	**bcsstk30**	28924	1007284
randomA4	1000	8177	**144**	144649	1074393
bcspwr10	5300	8271	**ct20stif**	52329	1273983
bcsstm13	649	9949	**m14b**	214765	1679018
dwt2680	2680	11173	**mrng2**	1017253	2015714
airfoil1	4253	12289	**auto**	448695	3314611
bcsstk12	1423	16342	**pwtk**	217918	5653257

Table 2: Results (small graphs).

Graph	**LB**	$\mathbf{\Delta_{Quick}}$	**ML**	**ML+GS**	**Quick**	**SP**	**+GS**	**+WM**	**+SA**	**HSA**
gd96c	1.25E+03	176.9	0.8	0.3	3.45500E+03	46.0	5.8	0.0	0.0	0.0
gd95c	1.13E+03	232.9	0.8	0.0	3.75500E+03	26.2	0.3	0.0	0.0	0.0
gd96b	2.43E+03	684.8	8.7	0.0	1.90860E+04	53.7	1.3	1.1	1.0	1.0
gd96d	3.74E+04	46.3	8.2	1.0	5.47390E+04	87.0	1.0	0.2	0.1	0.0
dwt245	4.34E+04	45.8	2.9	0.4	6.32810E+04	80.4	2.6	0.8	0.8	0.0
bintree10	8.85E+04	53.2	11.2	0.0	1.35656E+05	16394.2	38.8	11.3	10.3	6.4
bus685	1.43E+05	50.8	9.6	0.2	2.15744E+05	44.9	9.8	7.0	7.0	6.6
bus1138	4.00E+05	38.0	6.7	0.4	5.52111E+05	76.5	7.4	1.8	0.9	0.4
gd96a	2.83E+06	430.0	13.6	0.1	1.49741E+07	124.4	31.3	25.2	21.9	15.9
can445	1.57E+06	5.0	1.2	0.0	1.65431E+06	6.0	0.8	0.6	0.6	0.6
c1y	5.56E+06	41.4	8.2	0.0	7.86685E+06	121.1	5.6	4.6	4.3	4.2
c2y	8.74E+06	22.7	7.3	0.0	1.07286E+07	61.2	0.9	0.6	0.3	0.3
bcspwr08	7.97E+05	17.8	5.6	0.4	9.39437E+05	48.6	13.0	11.2	9.1	1.9
bcspwr09	7.91E+05	28.7	5.8	0.5	1.01801E+06	71.8	25.5	22.5	22.0	11.3
c5y	1.16E+07	21.2	7.2	0.0	1.39958E+07	130.5	10.8	9.6	8.4	6.4
jagmesh1	8.27E+05	5.1	2.5	0.1	8.68459E+05	14.2	12.6	12.2	11.8	1.1
c3y	1.55E+07	27.0	7.0	0.0	1.97321E+07	142.7	7.5	2.5	0.8	0.7
c4y	1.44E+07	15.0	8.2	0.0	1.66028E+07	51.8	2.1	1.4	0.2	0.0
dwt918	5.46E+05	51.0	5.1	0.1	8.25233E+05	11.5	1.4	0.9	0.3	0.0
dwt1007	8.86E+05	15.9	1.4	0.0	1.02750E+06	4.3	2.2	1.9	1.7	0.0
(continued)										

Table 3: Results (small graphs) - continuation. The average results are calculated for all the 37 graphs of Tables 2 and 3.

Graph	**LB**	$\mathbf{\Delta_{Quick}}$	**ML**	**ML+GS**	**Quick**	**SP**	**+GS**	**+WM**	**+SA**	**HSA**
jagmesh9	1.10E+06	27.4	5.2	0.2	1.39541E+06	10.3	6.3	4.5	1.6	0.9
can838	7.02E+06	5.8	0.4	0.0	7.43012E+06	1.8	0.1	0.1	0.0	0.0
randomA1	7.03E+07	321.8	34.9	1.8	2.96618E+08	49.7	18.2	9.7	4.9	1.1
hc10	1.79E+08	0.0	3.6	0.0	1.78957E+08	0.8	0.1	0.1	0.0	0.0
can1054	5.79E+06	9.9	0.2	0.1	6.36257E+06	2.0	0.1	0.1	0.0	0.0
can1072	8.17E+06	6.5	3.6	0.0	8.70400E+06	3.6	0.1	0.0	0.0	0.0
randomG4	7.33E+06	5.0	6.9	0.0	7.70221E+06	7.8	1.1	0.8	0.6	0.1
randomA4	3.01E+08	125.6	18.3	4.3	6.78008E+08	31.2	15.3	5.7	0.9	0.4
bcspwr10	1.19E+07	15.0	10.5	0.2	1.37238E+07	19.0	4.9	4.1	3.0	2.3
bcsstm13	2.23E+07	77.2	0.5	0.0	3.94573E+07	31.7	0.8	0.6	0.3	0.0
dwt2680	7.31E+06	25.6	4.2	0.0	9.18901E+06	5.2	0.3	0.1	0.0	0.0
airfoil1	1.18E+07	37.9	8.9	0.1	1.63343E+07	18.4	7.2	5.9	2.7	1.1
bcsstk12	1.71E+07	20.6	7.7	0.1	2.06281E+07	19.2	11.9	10.2	6.8	5.9
nasa1824	1.37E+08	3.1	5.8	0.0	1.41216E+08	24.0	7.9	4.2	1.1	0.3
randomA2	2.21E+09	33.5	12.8	4.1	2.95112E+09	12.9	5.1	0.3	0.1	0.0
nasa2146	1.11E+08	11.3	5.3	0.1	1.23584E+08	6.6	4.3	4.2	4.0	2.1
bcsstk13	4.35E+08	54.4	3.4	0.0	6.71461E+08	40.2	14.8	9.3	3.0	2.7
AVERAGE		75.4	6.9	0.4		41.3	7.5	4.7	3.5	2.0

Table 4: Results (large graphs).

Graph	**Quick**	**Graph**	**Quick**
whitaker3	6.53774E+07	**msc23052**	6.58277E+10
zcrack	1.36390E+08	**bcsstk36**	6.58053E+10
shuttleeddy	1.36200E+08	**bcsstk31**	7.45410E+10
randomA3	6.63612E+09	**msc10848**	5.95150E+10
nasa4704	7.54695E+08	**ferotor**	2.67776E+11
bcsstk24	9.06089E+08	**bcsstk35**	7.51880E+10
bcsstk38	3.87606E+09	**598a**	3.85388E+11
finan512	1.00967E+10	**bcsstk32**	1.46284E+11
bcsstk33	2.97010E+10	**bcsstk30**	5.11256E+10
bcsstk29	1.06444E+10	**144**	1.55347E+12
ocean	1.16999E+11	**ct20stif**	6.77425E+11
tooth	3.38761E+11	**m14b**	1.67209E+12
mrng1	6.69398E+11	**mrng2**	1.93775E+13
bcsstk37	6.77934E+10	**auto**	1.33598E+13
		pwtk	2.25527E+12

Table 5: The parameters used for the "**Quick**" and "**Extended**" V-cycles.

Parameter	"**Quick**" V-cycle	"**Extended**" V-cycle	The increase for level L
The coarse neighborhood size (r)	10	10	$+log(R)$
The edge filtering threshold (ϵ)	0.001	0.001	$\cdot 0.9^{log(R)}$
Compatible relaxation sweeps (k_1)	5	10	$+2 \cdot L$
GS relaxation sweeps (k_2)	5	10	$+2 \cdot L$
Heating and cooling in SA (k_3)	0	3	$\cdot log(R)$
k_4 used in the SA	0	5	$+log(\sqrt{R})$

Table 6: Running time

Graph	T_{Quick}	T_{Post}	Graph	T_{Quick}	T_{Post}
gd96c	*	*	**nasa1824**	0.05	0.17
gd95c	*	*	**randomA2**	0.26	0.27
gd96b	*	*	**nasa2146**	0.1	0.38
gd96d	*	*	**bcsstk13**	0.1	0.44
dwt245	*	0.01	**whitaker3**	0.15	-
bintree10	*	0.02	**zcrack**	0.13	-
bus685	*	0.02	**shuttleeddy**	0.15	-
bus1138	0.05	0.04	**randomA3**	0.42	-
gd96a	0.04	0.05	**nasa4704**	0.11	-
can445	*	0.02	**bcsstk24**	0.14	-
c1y	*	0.02	**bcsstk38**	0.3	-
c2y	0.02	0.03	**finan512**	1.2	-
bcspwr08	0.04	0.05	**bcsstk33**	0.65	-
bcspwr09	0.05	0.06	**bcsstk29**	0.6	-
c5y	0.02	0.06	**ocean**	7.5	-
jagmesh1	0.04	0.04	**tooth**	2.5	-
c3y	0.02	0.05	**mrng1**	12.3	-
c4y	0.02	0.05	**bcsstk37**	1.3	-
dwt918	*	0.04	**msc23052**	1.25	-
dwt1007	*	0.04	**bcsstk36**	1.3	-
jagmesh9	0.05	0.07	**bcsstk31**	1.9	-
can838	*	0.04	**msc10848**	1.1	-
randomA1	0.1	0.11	**ferotor**	4.7	-
hc10	0.065	0.06	**bcsstk35**	1.6	-
can1054	*	0.07	**598a**	6.4	-
can1072	*	0.05	**bcsstk32**	2.9	-
randomG4	0.02	0.07	**bcsstk30**	2	-
randomA4	0.15	0.18	**144**	10.3	-
bcspwr10	0.09	0.17	**ct20stif**	4.7	-
bcsstm13	0.04	0.14	**m14b**	17.5	-
dwt2680	0.05	0.12	**mrng2**	143	-
airfoil1	0.06	0.19	**auto**	64.3	-
bcsstk12	0.05	0.17	**pwtk**	20	-

References

[1] S. Barnard, A. Pothen, and H. Simon. A spectral algorithm for envelope reduction of sparse matrices. *Numerical Linear Algebra with Applications*, 2(4):317–334, 1995.

[2] A. Brandt. Algebraic multigrid theory: The symmetric case. 19:23–56, 1986. Preliminary proceedings of the International Multigrid Conference, April 6–8, 1983, Copper Mountain, CO.

[3] A. Brandt. General highly accurate algebraic coarsening. *Electronic Trans. Num. Anal. 10 (2000) 1-20*, 2000.

[4] A. Brandt. Multiscale scientific computation: Review 2001. 2001.

[5] A. Brandt, S. McCormick, and J. Ruge. Algebraic multigrid (AMG) for automatic multigrid solution with application to geodetic computations. Technical report, Institute for Computational Studies, Fort Collins, CO, POB 1852, 1982.

[6] A. Brandt, S. McCormick, and J. Ruge. Algebraic multigrid (AMG) for sparse matrix equations. In D. J. Evans, editor, *Sparsity and its Applications*, pages 257–284, 1984.

[7] A. Brandt and D. Ron. Chapter 1 : Multigrid solvers and multilevel optimization strategies. In J. Cong and J. R. Shinnerl, editors, *Multilevel Optimization and VLSICAD*. Kluwer, 2002.

[8] A. Brandt, D. Ron, and D. Amit. Multi-level approaches to discrete-state and stochastic problems. In W. Hackbush and U. Trottenberg, editors, *Multigrid Methods II*, pages 66–99. Springer-Verlag, 1986.

[9] W. L. Briggs, V. E. Henson, and S. F. McCormick. *A multigrid tutorial: second edition*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.

[10] C.-K. Cheng. Linear placement algorithm and applications to vlsi design. *Netw.*, 17(4):439–464, 1987.

[11] S. W.-S. D. Ron and A. Brandt. An algebraic multigrid based algorithm for bisectioning general graphs, 2005.

[12] T. Davis. University of florida sparse matrix collection. *NA Digest*, 97(23), 1997.

[13] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified np-complete problems. In *STOC '74: Proceedings of the sixth annual ACM symposium on Theory of computing*, pages 47–63, New York, NY, USA, 1974. ACM Press.

[14] A. George and A. Pothen. An analysis of spectral envelope reduction via quadratic assignment problems. *SIAM Journal on Matrix Analysis and Applications*, 18(3):706–732, 1997.

[15] K. Hall. An $r-$dimensional Quadratic Placement Algorithm. *Management Science*, 17:217–229, 1970.

[16] B. Hendrickson and R. Leland. An improved spectral graph partitioning algorithm for mapping parallel computations. *SIAM Journal on Scientific Computing*, 16(2):452–469, 1995.

[17] G. Karypis and V. Kumar. hmetis 1.5: A hypergraph partitioning package. *Tech. Report, Dept. of Computer Science, Univ. of Minnesota*, 1998.

[18] S. Kirkpatrick. Models of disordered systems. *Lecture Notes in Physics 149*.

[19] Y. Koren and D. Harel. Multi-scale algorithm for the linear arrangement problem. *Proceedings of 28th Inter. Workshop on Graph-Theoretic Concepts*, 2002.

[20] K. Mehlhorn and S. Näher. Leda: a platform for combinatorial and geometric computing. *Commun. ACM*, 38(1):96–102, 1995.

[21] J. Petit. Approximation heuristics and benchmarkings for the minla problem, 1998.

[22] A. Pothen, H. Simon, and L. Wang. Spectral nested dissection. Technical Report CS-92-01, 1992.

[23] A. Pothen, H. D. Simon, and K.-P. Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Matrix Anal. Appl.*, 11(3):430–452, 1990.

[24] R. Pozo, J. J. Dongarra, and D. W. Walker. Lapack++: a design overview of object-oriented extensions for high performance linear algebra. In *Supercomputing '93: Proceedings of the 1993 ACM/IEEE conference on Supercomputing*, pages 162–171, New York, NY, USA, 1993. ACM Press.

[25] J. Ruge and K. Stüben. *Algebraic Multigrid*, pages 73–130. SIAM, 1987.

[26] I. Safro, D. Ron, and A. Brandt. Graph minimum linear arrangement by multilevel weighted edge contractions. *Journal of Algorithms*, 2004, in press.

[27] I. Safro, D. Ron, and A. Brandt. Multilevel algorithms for linear ordering problems. *Submitted*, 2005.

[28] E. Sharon, A. Brandt, and R. Basri. Fast multiscale image segmentation. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, pages 70–77, 2000.

[29] D. A. Spielman and S.-H. Teng. Spectral partitioning works: Planar graphs and finite element meshes. In *IEEE Symposium on Foundations of Computer Science*, pages 96–105, 1996.

[30] K. Stüben. *An introduction to algebraic multigrid*, pages 413–532. Academic Press, 2001.

[31] K. Stüben. A review of algebraic multigrid. *J. Comput. Appl. Math.*, 128(1-2):281–309, 2001.