

ON LARGE SCALE DIAGONALIZATION TECHNIQUES FOR THE ANDERSON MODEL OF LOCALIZATION*

OLAF SCHENK[†], MATTHIAS BOLLHÖFER[‡], AND RUDOLF A. RÖMER[§]

Abstract. We propose efficient preconditioning algorithms for an eigenvalue problem arising in quantum physics, namely the computation of a few interior eigenvalues and their associated eigenvectors for the largest sparse real and symmetric indefinite matrices of the Anderson model of localization. We compare the Lanczos algorithm in the 1987 implementation by Cullum and Willoughby with the shift-and-invert techniques in the implicitly restarted Lanczos method and in the Jacobi-Davidson method. Our preconditioning approaches for the shift-and-invert symmetric indefinite linear system are based on maximum weighted matchings and algebraic multilevel incomplete LDL^T factorizations. These techniques can be seen as a complement to the alternative idea of using more complete pivoting techniques for the highly ill-conditioned symmetric indefinite Anderson matrices. We demonstrate the effectiveness and the numerical accuracy of these algorithms. Our numerical examples reveal that recent algebraic multilevel preconditioning solvers can accelerate the computation of a large-scale eigenvalue problem corresponding to the Anderson model of localization by several orders of magnitude.

Key words. Anderson model of localization, large-scale eigenvalue problem, Lanczos algorithm, Jacobi-Davidson algorithm, Cullum-Willoughby implementation, symmetric indefinite matrix, multilevel-preconditioning, maximum weighted matching

AMS subject classifications. 65F15, 65F50, 82B44, 65F10, 65F05, 05C85

1. Introduction. One of the hardest challenges in modern eigenvalue computation is the numerical solution of large-scale eigenvalue problems, in particular those arising from quantum physics such as, e.g., the Anderson model of localization (see Section 3 for details). Typically, these problems require the computation of some eigenvalues and -vectors for systems which have up to several million unknowns due to their high spatial dimensions. Furthermore, their underlying structure involves random perturbations of matrix elements which invalidates simple preconditioning approaches based on the graph of the matrices. Moreover, one is often interested in finding some eigenvalues and associated eigenvectors in the interior of the spectrum. The classical Lanczos approach [51] has led to eigenvalue algorithms [16, 17] that are in principle able to compute these eigenvalues using only a small amount of memory. More recent work on implicitly started Lanczos techniques [42, 43] has accelerated these methods significantly, yet to be fast one needs to combine this approach with shift-and-invert techniques, i. e. in every step one has to solve a shifted system of type $A - \sigma I$, where σ is a shift near the desired eigenvalues and $A \in \mathbb{R}^{n,n}$, $A = A^T$ is the associated matrix. In general shift-and-invert techniques converge rather quickly which is inline with the theory [51]. Still, an efficient solver is required to solve systems $(A - \sigma I)x = b$ efficiently with respect to time and memory. While implicitly restarted Lanczos techniques [42, 43] usually require to solve the system $(A - \sigma I)x = b$ to maximum precision and thus are mainly suited for sparse direct

*This work was supported by the Swiss Commission for Technology and Innovation (CTI) under grant number 7036 ENS-ES.

[†]Department of Computer Science Department, University Basel, Klingelbergstrasse 50, CH-4056 Basel, Switzerland, (olaf.schenk@unibas.ch).

[‡]Department of Mathematics, MA 4-5, TU Berlin, Str. des 17. Juni, 10623 Berlin, Germany, (bolle@math.tu-berlin.de). Supported by the DFG research center MATHEON “Mathematics for Key Technologies” in Berlin.

[§]Centre for Scientific Computing and Department of Physics, University of Warwick, Coventry CV4 9BU UK, (r.roemer@warwick.ac.uk).

solvers, the JACOBI-DAVIDSON method has become an attractive alternative [61] in particular when dealing with preconditioning methods for linear systems.

Until recently, sparse symmetric indefinite direct solvers were still far off from symmetric positive definite solvers and this might have been one major reason why shift-and-invert techniques were not able to compete with traditional Lanczos techniques [27], in particular because of memory constraints. With the invention of fast matchings-based algorithms [49] which improve the diagonal dominance of linear systems the situation has dramatically changed and the impact on preconditioning methods [7] as well as the benefits for sparse direct solvers [58] has been recognized. Furthermore, these techniques have been successfully transferred to the symmetric case [22, 24] allowing modern state-of-the-art direct solvers [57] to be orders of magnitudes faster and more memory efficient than ever, finally leading to symmetric indefinite sparse direct solvers that are almost as efficient as their symmetric positive definite counter parts. Recently this approach has also been utilized to construct incomplete factorizations [38] with similarly dramatic success. For a detailed survey on preconditioning techniques for large symmetric indefinite linear systems the interested reader should consult [5, 6].

2. Numerical approach for large systems. In the present paper we combine the above mentioned advances with inverse-based preconditioning techniques [8]. This allows us to find interior eigenvalues and -vectors for the Anderson problem several orders of magnitudes faster than traditional algorithms [16, 17] while still keeping the amount of memory reasonably small.

Let us briefly outline our strategy. We will consider recent novel approaches in preconditioning methods for symmetric indefinite linear systems and eigenvalue problems and apply them to the Anderson model. Since the Anderson model is a large-scale sparse eigenvalue problem in three spatial dimensions, the eigenvalue solvers we deal with are designed to compute only a few interior eigenvalues and eigenvectors avoiding a complete factorization. In particular we will use two modern eigenvalue solvers which we will briefly introduce in Section 4. The first one is ARPACK [42], which is a Lanczos-type method using implicit restarts (cf. section 4.1). We use this algorithm together with a shift-and-invert technique, i. e. eigenvalues and eigenvectors of $(A - \sigma I)^{-1}$ are computed instead of those of A . ARPACK is used in conjunction with a direct factorization method and a multilevel incomplete factorization method for the shift-and-invert technique.

Firstly, we use the shift-and-invert technique with the novel symmetric indefinite sparse direct solver that is part of PARDISO [57] and we report extensive numerical results on the performance of this method. Section 5 will give a short overview of the main concepts that form the PARDISO solver. Secondly, we use ARPACK in combination with the multilevel incomplete LU factorization package ILUPACK [9]. Here we present a new *indefinite* version of this preconditioner that is devoted to symmetrically indefinite problems and combines two basic ideas, namely (i) symmetric maximum weighted matchings [22, 24] and (ii) inverse-based decomposition techniques [8]. These will be described in Sections 5.2 and 7.

As a second eigenvalue solver we use the symmetric version of the JACOBI-DAVIDSON method, in particular the implementation JDBSYM [32]. This Newton-type method (see section 4.2) is used together with ILUPACK [9]. As we will see in several further numerical experiments, the synergy of both approaches will form an extremely efficient preconditioner for the Anderson model that is memory efficient while at the same time accelerates the eigenvalue computations significantly: system sizes that

resulted in weeks of computing time [27] can now be computed within an hour.

3. The Anderson model of localization. The Anderson model of localization is a paradigmatic model describing the electronic transport properties of disordered quantum systems [41, 54]. It has been used successfully in amorphous materials such as alloys [52], semiconductors and even DNA [53]. Its hallmark is the prediction of a spatial confinement of the electronic motion upon increasing the disorder — the so-called Anderson *localization* [2]. When the model is used in 3 spatial dimensions, it exhibits a metal-insulator transition in which the disorder strength w mediates a change of transport properties from metallic behavior at small w via critical behavior at the transition w_c and on to insulating behavior and strong localization at larger w [41, 54]. Mathematically, the quantum problem corresponds to a Hamilton operator in the form of a real symmetric matrix A , with quantum mechanical energy levels given by the eigenvalues $\{\lambda\}$, and the respective wave functions are simply the eigenvectors of A , i.e. vectors x with real entries. With $N = M \times M \times M$ sites, the quantum mechanical (stationary) Schrödinger equation is equivalent to the eigenvalue equation $Ax = \lambda x$, which in site representation reads as

$$x_{i-1;j;k} + x_{i+1;j;k} + x_{i;j-1;k} + x_{i;j+1;k} + x_{i;j;k-1} + x_{i;j;k+1} + \varepsilon_{i;j;k}x_{i;j;k} = \lambda x_{i;j;k} \quad (3.1)$$

with $i, j, k = 1, 2, \dots, M$ denoting the Cartesian coordinates of a site. The disorder enters the matrix on the diagonal where the entries $\varepsilon_{i;j;k}$ correspond to a spatially varying disorder potential and are selected randomly according to a suitable distribution [40]. Here, we shall use the standard box distribution $\varepsilon_{i;j;k} \in [-w/2, w/2]$ such that the w parameterizes the aforementioned disorder strength. Clearly, the eigenvalues of A then lie within the interval $[-6-w/2, 6+w/2]$. In most studies of the induced metal-insulator transition, w ranges from 1 to 30 [54]. But these values also depend on whether generalizations to random off-diagonal elements [26, 63] — the so-called random-hopping problem —, anisotropies [45, 48] or other lattice graphs [36, 60] are being considered.

The intrinsic physics of the model is quite rich. For disorders $w \ll 16.5$, the eigenvectors are extended, i.e. $x_{i;j;k}$ is fluctuating from site to site, but the envelope $|x|$ is approximately a nonzero constant. For large disorders $w > 16.5$, all eigenvectors are localized such that the envelope $|x_n|$ of the n th eigenstate may be approximately written as $\exp -[|\vec{r} - \vec{r}_n|]/l_n(w)$ with $\vec{r} = (i, j, k)^T$ and $l_n(w)$ denoting the *localization length* of the eigenstate. In Figure 3.1, we show examples of such states. Note that $|x|^2$ and not x corresponds to a physically measurable quantity and is therefore the observable of interest to physicists. Directly at $w = w_c \approx 16.5$, the extended states at $\lambda = 0$ vanish and no current can flow. The wave function vector x appears simultaneously extended and localized as shown in Fig. 3.2.

In order to numerically distinguish these three regimes, namely, localized, critical, and extended behaviors, one needs to (i) go to extremely large system sizes of order 10^6 to 10^8 and (ii) average over many different realizations of the disorder, i.e., compute eigenvalues or eigenvectors for many matrices with different diagonals. In the present paper, we concentrate on the computation of a few eigenvalues and corresponding eigenvectors for the physically most interesting case of critical disorder w_c and in the center of $\sigma(A)$, i.e., at $\lambda = 0$, for large system sizes [3, 10, 47, 64]. Since there is a high density of states for $\sigma(A)$ at $\lambda = 0$ in all cases, we have the further numerical challenge of clearly distinguishing the eigenstates in this high density region.

3.1. Lanczos algorithm and the Cullum-Willoughby implementation.

Since the mid-eighties, the preferred numerical tool to study the Anderson matrix and

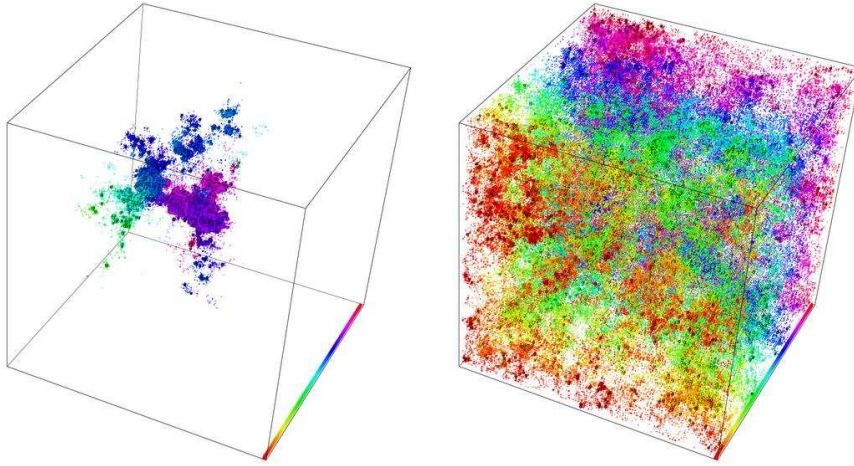


FIG. 3.1. *Extended (left) and localized (right) wave function probabilities for the 3D Anderson model with periodic boundary conditions at $\lambda =$ with $N = 100^3$ and $w = 12.0$ and 21.0 , respectively. Every site with probability $|x_j|^2$ larger than the average $1/N^3$ is shown as a box with volume $|x_j|^2 N$. Boxes with $|x_j|^2 N > \sqrt{1000}$ are plotted with black edges. The color scale distinguishes between different slices of the system along the axis into the page.*

to compute a selected set of eigenvectors, e.g. as needed for a multifractal analysis at the transition, was provided by the Cullum-Willoughby implementation (CWI) [16–18] of the Lanczos algorithm. The algorithm iteratively generates a sequence of orthogonal vectors v_i , $i = 1, \dots, K$, such that $V_K^T A V_K = T_K$, with $V = [v_1, v_2, \dots, v_K]$ and T_K a symmetric tridiagonal $K \times K$ matrix. In exact arithmetic, the recursion

$$\beta_{i+1} v_{i+1} = A v_i - \alpha_i v_i - \beta_i v_{i-1}, \quad (3.2)$$

where $\alpha_i = v_i^T A v_i$ and $\beta_{i+1} = v_{i+1}^T A v_i$ are the diagonal and subdiagonal entries of T_K and $v_0 = 0$ and v_1 is an arbitrary starting vector, is an orthogonal transformation to tridiagonal form that needs $K = N$ matrix-vector multiplications. The eigenvalues of the tridiagonal matrix T_K (Ritz values) are then simply the eigenvalues of the matrix A and the associated Ritz vectors yield the eigenvectors.

Since A is sparse and symmetric, the underlying matrix-vector multiplication on the CWI can be programmed very efficiently, either by directly coding or appropriate sparse storage schemes — only the diagonal needs to be stored in any case. Additionally, the CWI is a Lanczos implementation in which *no reorthogonalization* is performed. Rather, spurious eigenvalues are identified by extending the set of Ritz vectors to more than the N present in exact arithmetic. Typically, we find that $K \approx 4N$ is sufficient for all “good” eigenvalues to have replicated themselves at least twice — a further sign that the algorithm has converged. Clearly, this strategy work well in the present case since the disorder destroys any symmetry-induced degeneracies.

The CWI is memory efficient and does not need elaborate reorthogonalization schemes, but does need to construct many Ritz vectors which is computationally intensive. Nevertheless, in 1999 CWI was still significantly faster than more modern iterative schemes [27]. The main reason for this surprising result lies in the indefiniteness of A , which led to severe difficulties with solvers more accustomed to standard




FIG. 3.2. *Plot of the electronic eigenstate at the metal-insulator transition with $E = 0$, $w = 16.5$ and $N = 250^3$. The box-and-color scheme is as in Fig. 3.1. Note how the state extends nearly everywhere while at the same time exhibiting certain localized regions of higher $|x_j|^2$ values.*

Laplacian-type problems.

4. Modern approaches for solving symmetric indefinite eigenvalue problems. When dealing with eigenvalues near a given real shift σ , the Lanczos algorithm [51] is usually accelerated when being applied to the shifted inverse $(A - \sigma I)^{-1}$ instead of A directly. This approach relies on the availability of a fast solution method for linear systems of type $(A - \sigma I)x = b$. However, the limited amount of available memory only allows for a small number of solution steps and sparse direct solvers also need to be memory-efficient to turn this approach into a practical method.

The limited number of Lanczos steps has led to modern implicitly restarted methods [43, 62] which ensure that the information about the desired eigenvalues is inherited when being restarted. With increasing number of preconditioned itera-

tive methods for linear systems [55], Lanczos-type algorithms have become less attractive mainly because in every iteration step the systems of type $(A - \sigma I)x = b$ have to be solved to *full* accuracy in order to avoid false eigenvalues. In contrast to this, JACOBI-DAVIDSON-like methods [61] allow using a crude approximation of the underlying linear system. From the point of view of linear solvers as part of the eigenvalue computation, modern direct and iterative methods need to inherit the symmetric structure $A = A^T$ while remaining both time and memory efficient. Symmetric matching algorithms [22, 24, 57] have significantly improved these methods.

4.1. The shift-and-invert mode of the restarted Lanczos method. The Lanczos method for real symmetric matrices A near a shift σ is based on computing successively orthonormal vectors $[v_1, \dots, v_k, v_{k+1}]$ and a tridiagonal $(k+1) \times k$ matrix

$$\underline{T}_k = \begin{pmatrix} \alpha_1 & \beta_1 & & & \\ \beta_1 & \alpha_2 & & & \\ & & \ddots & & \\ & & & \ddots & \beta_{k-1} \\ & & & \beta_{k-1} & \alpha_k \\ \hline & & & & & \beta_k \end{pmatrix} \equiv \begin{pmatrix} T_k \\ \beta_k e_k^T \end{pmatrix}, \quad (4.1)$$

where e_k is the k th unit vector in \mathbb{R}^k , such that

$$(A - \sigma I)^{-1}[v_1, \dots, v_k] = [v_1, \dots, v_k, v_{k+1}]\underline{T}_k. \quad (4.2)$$

Since only a limited number of Lanczos vectors v_1, \dots, v_k can be stored and since this Lanczos sequence also consists of redundant information about undesired small eigenvalues, implicitly restarted Lanczos methods have been proposed [43, 62] that use implicitly shifted QR [35] exploiting the small eigenvalues of T_k to remove them out of this sequence without ever forming a single matrix vector multiplication with $(A - \sigma I)^{-1}$. The new transformed Lanczos sequence

$$(A - \sigma I)^{-1}[\tilde{v}_1, \dots, \tilde{v}_l] = [\tilde{v}_1, \dots, \tilde{v}_l, \tilde{v}_{l+1}]\underline{\tilde{T}}_l \quad (4.3)$$

with $l \ll k$ then allows to compute further $k - l$ approximations. This approach is at the heart of the symmetric version of ARPACK [42, 43].

4.2. The symmetric Jacobi-Davidson method. One of the major drawbacks of shift-and-invert Lanczos algorithms is the fact that the multiplication with $(A - \sigma I)^{-1}$ requires solving a linear system to full accuracy. In contrast to this, JACOBI-DAVIDSON-like algorithms [61] are based on a Newton-like approach to solve the eigenvalue problem. Like the Lanczos method the search space is expanded step by step solving the correction equation

$$(I - uu^T)(A - \theta I)(I - uu^T)z = -r \quad \text{such that} \quad z = (I - uu^T)z \quad (4.4)$$

where (u, θ) is the given approximate eigenpair and $r = Au - \theta u$ is the associated residual. Then the search space based on $V_k = [v_1, \dots, v_k]$ is expanded by reorthogonalizing z with respect to $[v_1, \dots, v_k]$ and a new approximate eigenpair is computed from the Ritz approximation $[V_k, z]^T A [V_k, z]$. When computing several right eigenvectors, the projection $I - uu^T$ has to be replaced by $I - [Q, u][Q, u]^T$ using the already computed approximate eigenvectors Q . This ensures that the new approximate eigenpair is orthogonal to those that have already been computed.

The most important part of the JACOBI-DAVIDSON approach is to construct an approximate solution for (4.4) such that

$$(I - uu^T)K(I - uu^T)c = d \quad \text{with} \quad u^T z = 0 \quad (4.5)$$

and $K \approx A - \theta I$ that allows a fast solution of the system $Kx = b$. Here, there is a strong need for robust preconditioning methods that preserve symmetry and efficiently solve sequences of linear systems with K . If K is itself symmetric and indefinite, then the simplified QMR method [29, 30] using the preconditioner $(I - \frac{uw^T}{w^T u})K^{-1}$, where $Kw = u$ and the system matrix $(I - uu^T)(A - \theta I)$ can be used as iterative method. Note that here the accuracy of the solution of (4.4) is uncritical until the approximate eigenpair converges [28]. This fact has been exploited in JDBSYM [4, 32]. For an overview on JACOBI-DAVIDSON methods for symmetric matrices see [33].

5. On recent algorithms for solving symmetric indefinite systems of equations. We now report on recent improvements in solving symmetric indefinite systems of linear equations that have significantly changed sparse direct as well as preconditioning methods. One key role that made these approaches successful is played by the use of symmetric matchings that we review in Section 5.2.

5.1. Sparse direct factorization methods. For a long time dynamic pivoting has been a central point for nonsymmetric sparse linear solvers to gain stability. Therefore, improvements in speeding up direct factorization methods were limited to the uncertainties that have arisen from using pivoting. Certainly techniques like the column elimination tree [19, 34] have been a useful tool to predict the sparsity pattern despite pivoting. However, in the symmetric case the situation becomes more complicated since only symmetric reorderings applied to both, columns and rows, are required and no a-priori choice of pivots is given. This makes it almost impossible to predict the elimination tree in a sensible manner and the use of cache-oriented level-3 BLAS [20, 21] was impossible.

With the introduction of symmetric maximum weighted matchings [22] as alternative to complete pivoting it is now possible to treat symmetric indefinite systems almost similar to symmetric positive definite systems. This allows the prediction of fill using the elimination tree [31] and thus to set up the data structures that are required to predict dense submatrices (also known as supernodes). This in turn means that one is able to exploit level-3 BLAS applied to the supernodes. Consequently, the classical Bunch-Kaufman pivoting approach [12] need to be performed only inside the supernodes.

This approach has recently been successfully implemented in the sparse direct solver PARDISO [57] and as a major consequence, this novel approach has improved the sparse indefinite solver to become almost as efficient as its symmetric positive analogy. Certainly for the Anderson problem studied here, PARDISO is about 2 orders of magnitude more efficient than previously used direct solvers [27]. We also note that the idea of symmetric weighted matchings can be carried over to incomplete factorization methods with similar success [38].

5.2. Symmetric weighted matchings as an alternative to complete pivoting techniques. Symmetric weighted matchings [22, 24], which will be explained in detail in Section 6.2, can be viewed as a preprocessing step that rescales the original matrix and at the same time improves the block diagonal dominance. By this strategy, all entries are at most one in modulus and in addition the diagonal blocks are

either 1×1 scalars a_{ii} such that $|a_{ii}| = 1$ (in exceptional cases we will have $a_{ii} = 0$) or they are 2×2 blocks

$$\begin{pmatrix} a_{ii} & a_{i,i+1} \\ a_{i+1,i} & a_{i+1,i+1} \end{pmatrix} \text{ such that } |a_{ii}|, |a_{i+1,i+1}| \leq 1 \text{ and } |a_{i+1,i}| = |a_{i,i+1}| = 1.$$

Although this strategy does not necessarily ensure that symmetric pivoting like in Ref. [12] is unnecessary, it is nevertheless likely to waive dynamic pivoting during the factorization process. It has been shown in [24] that based on symmetric weighted matchings the performance of the sparse symmetric indefinite *multifrontal* direct solver MA57 is improved significantly, although a dynamic pivoting strategy by Duff and Reid [25] was still present. Recent results in [57] have shown that the absence of dynamic pivoting does not harm the method anymore and that therefore symmetric weighted matchings can be considered as alternative to complete pivoting.

6. Symmetric reorderings to improve the results of pivoting on restricted subsets. In this section we will discuss weighted graph matchings as an additional preprocessing step. The motivation for weighted matching approaches is to identify large entries in the coefficient matrix A that, if permuted close to the diagonal, permit the factorization process to identify more acceptable pivots and proceed with fewer pivot perturbations. These methods are based on maximum weighted matchings \mathcal{M} and improve the quality of the factor in a complementary way to the alternative idea of using more complete pivoting techniques. The idea to use a permutation $P_{\mathcal{M}}$ associated with a weighted matching \mathcal{M} as an approximation of the pivoting order for nonsymmetric linear systems was firstly introduced by Olschowka and Neumaier [49] and extended by Duff and Koster [23] to the sparse case. Permuting the rows $A \leftarrow P_{\mathcal{M}}A$ of the sparse system to ensure a zero-free diagonal or to maximize the product of the absolute values of the diagonal entries are techniques that are now often regularly used for nonsymmetric matrices [7, 46, 58, 59].

6.1. Matching algorithms for nonsymmetric matrices. Let $A = (a_{ij}) \in \mathbb{R}^{n \times n}$ be a general matrix. The nonzero elements of A define a graph with edges $\mathcal{E} = \{(i, j) : a_{ij} \neq 0\}$ of ordered pairs of row and column indices. A subset $\mathcal{M} \subset \mathcal{E}$ is called a matching or a transversal, if every row index i and every column index j appear at most once in \mathcal{M} . A matching \mathcal{M} is called perfect if its cardinality is n . For a nonsingular matrix at least one perfect matching exists and can be found with well known algorithms. With a perfect matching \mathcal{M} , it is possible to define a permutation matrix $P_{\mathcal{M}} = (p_{ij})$ with:

$$p_{ij} = \begin{cases} 1 & (j, i) \in \mathcal{M}, \\ 0 & \text{otherwise.} \end{cases} \quad (6.1)$$

As a consequence, the permutation matrix $P_{\mathcal{M}}A$ has nonzero elements on its diagonal. This method only takes the nonzero structure of the matrix into account. There are other approaches which maximize the diagonal values in some sense. One possibility is to look for a matrix $P_{\mathcal{M}}$, such that the product of the diagonal values of $P_{\mathcal{M}}A$ is maximal. In other words, a permutation σ has to be found, which maximizes

$$\prod_{i=1}^n |a_{\sigma(i)i}|. \quad (6.2)$$




FIG. 6.1. *Illustration of the row permutation. A small numerical value is indicated by a \circ -symbol and a large numerical value by an \bullet -symbol. The matched entries \mathcal{M} are marked with squares and $P_{\mathcal{M}} = (e_4; e_1; e_5; e_2; e_3; e_6)$.*

This maximization problem is solved indirectly. It can be reformulated by defining a matrix $C = (c_{ij})$ with

$$c_{ij} = \begin{cases} \log a_i - \log |a_{ij}| & a_{ij} \neq 0 \\ \infty & \text{otherwise,} \end{cases} \quad (6.3)$$

where $a_i = \max_j |a_{ij}|$, i.e. the maximum element in row i of matrix A . A permutation σ , which minimizes $\sum_{i=1}^n c_{\sigma(i)i}$ also maximizes the product (6.2).

The minimization problem is known as linear sum assignment problem or bipartite weighted matching problem in combinatorial optimization. The problem is solved by a sparse variant of the Kuhn-Munkres algorithm. The complexity is $O(n^3)$ for full $n \times n$ matrices and $O(n\tau \log n)$ for sparse matrices with τ entries. For matrices whose associated graph fulfills special requirements, this bound can be reduced further to $O(n^\alpha(\tau + n \log n))$ with $\alpha < 1$. All graphs arising from finite-difference or finite-element discretizations meet these conditions [37]. As before, we finally get a perfect matching \mathcal{M} that in turn defines a nonsymmetric permutation $P_{\mathcal{M}}$.

The effect of nonsymmetric row permutations using a permutation associated with a matching \mathcal{M} is shown in Figure 6.1. It is clearly visible that the matrix $P_{\mathcal{M}}A$ is now nonsymmetric, but has the largest nonzeros on the diagonal.

6.2. Symmetric 1×1 and 2×2 block weighted matchings. In the case of symmetric indefinite matrices, we are interested in symmetrically permuting PAP^T . The problem is that zero or small diagonal elements of A remain on the diagonal by using a symmetric permutation PAP^T . Alternatively, instead of permuting a large¹ off-diagonal element a_{ij} nonsymmetrically to the diagonal, one can try to devise a permutation P_S such that $P_S A P_S^T$ permutes this element close to the diagonal. As a result, if we form the corresponding 2×2 block to $\begin{bmatrix} a_{ii} & a_{ij} \\ a_{ij} & a_{jj} \end{bmatrix}$, we expect the off-diagonal entry a_{ij} to be large and thus the 2×2 block would form a suitable 2×2 pivot for the Supernode-Bunch-Kaufman factorization. An observation on how to build P_S from the information given by a weighted matching \mathcal{M} was presented by Duff and Gilbert [22]. They noticed that the cycle structure of the permutation $P_{\mathcal{M}}$ associated with the nonsymmetric matching \mathcal{M} can be exploited to derive such a permutation P_S . For example, the permutation $P_{\mathcal{M}}$ from Figure 6.1 can be written in cycle representation as $P_C = (e_1; e_2; e_4)(e_3; e_5)(e_6)$. This is shown in the upper graphics in Figure 6.2. The left graphic displays the cycles (1 2 4), (3 5) and (6). If we modify the original permutation $P_{\mathcal{M}} = (e_4; e_1; e_5; e_2; e_3; e_6)$ into this cycle permutation $P_C =$

¹Large in the sense of the weighted matching \mathcal{M} .




FIG. 6.2. Illustration of a cycle permutation with $P_C = (e_1; e_2; e_4)(e_3; e_5)(e_6)$ and $P_S = (e_1)(e_2; e_4)(e_3; e_5)(e_6)$. The symmetric matching P_S has two additional elements (indicated by dashed boxes), while one element of the original matching fell out (dotted box). The two 2-cycles are permuted into 2×2 diagonal blocks to serve as initial 2×2 pivots.

$(e_1; e_2; e_4)(e_3; e_5)(e_6)$ and permute A symmetrically with $P_C A P_C^T$, it can be observed that the largest elements are permuted to diagonal blocks. These diagonal blocks are shown by filled boxes in the upper right matrix. Unfortunately, a long cycle would result into a large diagonal block and the fill-in of the factor for $P_C A P_C^T$ may be prohibitively large. Therefore, long cycles corresponding to P_M must be broken down into disjoint 2×2 and 1×1 cycles. These smaller cycles are used to define a symmetric permutation $P_S = (c_1, \dots, c_m)$, where m is the total number of 2×2 and 1×1 cycles.

The rule for choosing the 2×2 and 1×1 cycles from P_C to build P_S is straightforward. One has to distinguish between cycles of even and odd length. It is always possible to break down even cycles into cycles of length two. For each even cycle, there are two possibilities to break it down. We use a structural metric [24] to decide which one to take. The same metric is also used for cycles of odd length, but the situation is slightly different. Cycles of length $2l + 1$ can be broken down into l cycles of length two and one cycle of length one. There are $2l + 1$ different possibilities to do this. The resulting 2×2 blocks will contain the matched elements of \mathcal{M} . However, there is no guarantee that the remaining diagonal element corresponding to the cycle of length one will be nonzero. Our implementation will randomly select one element as a 1×1 cycle from an odd cycle of length $2l + 1$.

A selection of P_S from a weighted matching P_M is illustrated in Figure 6.2. The permutation associated with the weighted matching, which is sorted according to the cycles, consists of $P_C = (e_1; e_2; e_4)(e_3; e_5)(e_6)$. We now split the full cycle of odd length three into two cycles $(1)(24)$ — resulting in $P_S = (e_1)(e_2; e_4)(e_3; e_5)(e_6)$. If P_S is symmetrically applied to $A \leftarrow P_S A P_S^T$, we see that the large elements from the nonsymmetric weighted matching \mathcal{M} will be permuted close to the diagonal and these elements will have more chances to form good initial 1×1 and 2×2 pivots for the subsequent (incomplete) factorization.

Good fill-in reducing orderings P_{Fill} are equally important for symmetric indefinite systems. The following section introduces two strategies to combine these reorderings with the symmetric graph matching permutation P_S . This will provide good initial pivots for the factorization as well as a good fill-in reduction permutation.

6.3. Combination of orderings P_{Fill} for fill reduction with orderings P_S based on weighted matchings. In order to construct the factorization efficiently, care has to be taken that not too much fill-in is introduced during the elimination process. We now examine two algorithms for the combination of a permutation P_S based on weighted matchings to improve the numerical quality of the coefficient matrix A with a fill-in reordering P_{Fill} based on a nested dissection from METIS [39]. The first method is based on compressed subgraphs and has also been used by Duff and Pralet in [24] in order to find good scalings and orderings for symmetric indefinite systems.

In order to combine the permutation P_S with a fill-in reducing permutation, we compress the graph of the reordered system $P_S A P_S^T$ and apply the fill-in reducing reordering to the compressed graph. In the compression step, the union of the structure of the two rows and columns corresponding to a 2×2 diagonal block are built, and used as the structure of a single, compressed row and column representing the original ones.

If $G_A = (V; E)$ is the undirected graph of A and a cycle consists of two vertices $(s, t) \in V$, then graph compression will be done on the 1×1 and 2×2 cycles, which have been found using a weighted matching \mathcal{M} on the graph. The vertices (s, t) are replaced with a single supervertex $u = \{s, t\} \in V_c$ in the compressed graph $G_c = (V_c, E_c)$. An edge $e_c = (s, t) \in E_c$ between two supervertices $s = \{s_1, s_2\} \in V_c$ and $t = \{t_1, t_2\} \in V_c$ exists if at least one of the following edges exist in E : $(s_1, t_1), (s_1, t_2), (s_2, t_1)$ or (s_2, t_2) . The fill-in reducing ordering is found by applying METIS on the compressed graph $G_c = (V_c, E_c)$. Expansion of that permutation to the original numbering yields P_{fill} . Hence all 2×2 cycles that correspond to a suitable 2×2 pivot block are reordered consecutively in the factor.

7. Symmetric multi-level preconditioning techniques. We now present a new symmetric indefinite approximate multilevel factorization that is mainly based on three parts which are repeated in a multilevel framework in each subsystem. The components consist of (i) reordering of the system, (ii) approximate factorization using inverse-based pivoting and, (iii) recursive application to the system of postponed updates.

7.1. Reordering the given system. The key ingredient to turn this approach into an efficient multilevel solver consists of the symmetric maximum weight matching presented in Section 5.2. After the system is reordered into a representation

$$P_s^T D A D P_s = \hat{A}, \tag{7.1}$$

where $D, P_s \in \mathbb{R}^{n,n}$, D is a diagonal matrix and P_s is a permutation matrix, \hat{A} is expected to have many diagonal blocks of size 1×1 or 2×2 that are well-conditioned. Once the diagonal block of size 1×1 and 2×2 are built, the associated block graph of \hat{A} is reordered by a symmetric reordering, e. g. AMD [1] or METIS [39], i. e.

$$\Pi^T P_s^T D A D P_s \Pi = \tilde{A}, \tag{7.2}$$

where $\Pi \in \mathbb{R}^{n,n}$ refers to the associated symmetric block permutation.

7.2. Inverse-based pivoting. Given \tilde{A} we compute an incomplete factorization $LDL^T = \tilde{A} + E$ of \tilde{A} . To do this at step k of the algorithm we have

$$\tilde{A} = \begin{pmatrix} B & F^T \\ F & C \end{pmatrix} = \begin{pmatrix} L_B & 0 \\ L_F & I \end{pmatrix} \begin{pmatrix} D_B & 0 \\ 0 & S_C \end{pmatrix} \begin{pmatrix} L_B^T & L_F^T \\ 0 & I \end{pmatrix}, \tag{7.3}$$

where $L_B \in \mathbb{R}^{k-1, k-1}$ is lower triangular with unit diagonal and $D_B \in \mathbb{R}^{k-1, k-1}$ is block diagonal with diagonal blocks of size 1×1 and 2×2 . Also, $S_C = C - L_F D_B L_F^T = (s_{ij})_{i,j}$ denotes the approximate Schur complement. To proceed with the incomplete factorization we perform either a 1×1 update or a 2×2 block update. One possible choice could be to use Bunch's algorithm [11]. This approach has been used in Ref. [38]. Here we use a simple criterion based on block diagonal dominance of the leading block column. Depending on the values

$$d_1 = \sum_{j>1} \frac{|s_{j1}|}{|s_{11}|}, \quad d_2 = \sum_{j>2} \|(s_{j1}, s_{j2}) \begin{pmatrix} s_{11} & s_{12} \\ s_{12} & s_{22} \end{pmatrix}^{-1}\|, \quad (7.4)$$

we perform a 2×2 update only if $d_2 < d_1$. The leading two columns of S_C can be efficiently computed using linked lists [44] and it is not required to have all entries of S_C available.

When applying the (incomplete) factorization LDL^T to \tilde{A} we may still encounter a situation where at step k either $1/|s_{11}|$ or $\|(s_{ij})_{i,j \leq 2}^{-1}\|$ is large or even infinity. Since we are dealing with an incomplete factorization we propose to use inverse-based pivoting [8]. Therefore we require in every step that

$$\left\| \begin{pmatrix} L_B & 0 \\ L_E & I \end{pmatrix}^{-1} \right\| \leq \kappa \quad (7.5)$$

for a prescribed bound κ . If after the update using a 1×1 pivot (or 2×2 pivot) the norm of the inverse lower triangular factor fails to be less than κ , the update is postponed and the leading rows/columns of L_E , S_C are permuted to the end. Otherwise depending on whether a 1×1 or a 2×2 pivot has been selected, the entries

$$(s_{j1}/s_{11})_{j>1}, \quad \left((s_{j1}, s_{j2}) \begin{pmatrix} s_{11} & s_{12} \\ s_{12} & s_{22} \end{pmatrix}^{-1} \right)_{j>2} \quad (7.6)$$

become the next (block) column of L and we drop these entries whenever their absolute value is less than ε/κ for some threshold ε . For a detailed description see Ref. [8]. The norm of the inverse can cheaply be estimated using a refined strategy of Ref. [15] and is part of the software package ILUPACK that is now extended to the symmetric indefinite case [9].

7.3. Recursive application. After the inverse-based ILU we have an approximate factorization

$$Q^T \tilde{A} Q = \begin{pmatrix} L_{11} & 0 \\ L_{21} & I \end{pmatrix} \begin{pmatrix} D_{11} & 0 \\ 0 & S_{22} \end{pmatrix} \begin{pmatrix} L_{11}^T & L_{21}^T \\ 0 & I \end{pmatrix} \quad (7.7)$$

and it typically does not pay off to continue the factorization for the remaining matrix S_{22} which consists of the previously postponed updates. Thus S_{22} is now explicitly computed and the strategies for reordering, scaling and factorization are recursively applied to S_{22} leading to a multilevel factorization.

Note that in order to save memory L_{21} is not stored but implicitly approximated by $\tilde{A}_{21}(L_{11}D_{11}U_{11})^{-1}$. In addition we use a technique called *aggressive dropping* that sparsifies the triangular factor L a posteriori. To do this observe that when applying a perturbed triangular factor \tilde{L}^{-1} for preconditioning instead of L^{-1} we have

$$\tilde{L}^{-1} = (I + E_L)L^{-1}, \quad \text{where } E_L = \tilde{L}^{-1}(L - \tilde{L}).$$

We can expect that \tilde{L}^{-1} serves as a good approximation to L^{-1} as long as $\|E_L\| \ll 1$. If we obtain \tilde{L} from L by dropping some entry, say l_{ij} from L , then we have to ensure that

$$\|\tilde{L}^{-1}e_i\| \cdot |l_{ij}| \leq \tau \ll 1,$$

for some moderate constant $\tau < 1$, e.g. $\tau = 0.1$. To do this it is required to have a good estimate for $\nu_i \approx \|\tilde{L}^{-1}e_i\|$ available, for any $i = 1, \dots, n$. In principle it can be computed [8, 15] using \tilde{L}^\top instead of \tilde{L} . Last, knowing how many entries exist in column j , we could drop any l_{ij} such that

$$|l_{ij}| \leq \tau / (\nu_i \cdot \#\{l_{kj} : l_{kj} \neq 0, k = j + 1, \dots, n\}).$$

7.4. Iterative solution. By construction, the computed incomplete multilevel factorization is symmetric but indefinite. For the iterative solution of linear systems using the multilevel factorization, in principle different Krylov subspace solvers could be used. For example, general methods that do not explicitly use symmetry (e.g. GMRES [56]) or methods like SYMMLQ [50] which preserve the symmetry of the original matrix but which are only devoted for symmetric positive definite preconditioners. To fully exploit both, symmetry and indefiniteness at the same time, here the simplified QMR method [29, 30] is chosen.

8. Numerical Experiments. Here we present numerical experiments that show that the previously outlined advances in symmetric indefinite sparse direct solvers as well as in preconditioning methods significantly accelerate modern eigenvalue solvers and allow us to gain orders of magnitude in speed compared to more conventional methods.

8.1. Computing Environments and Software. All large scale numerical experiments for the Anderson model of localization were performed on an SGI Altix 3700/BX2 with 56 Intel Itanium2 1.6 GHz processors and 112 GB of memory. If not explicitly stated, we always used only one processor of the system and all algorithms were implemented in either C or Fortran77. All codes were compiled by the Intel V8.1 compiler suite using *ifort* and *icc* with the -O3 optimization option and linked with basic linear algebra subprograms optimized for Intel architectures. For completeness, let us recall the main software packages used.

- ARPACK is a collection of Fortran77 subroutines designed to solve large scale eigenvalue problems. The eigenvalue solver has been developed at the Department of Computational and Applied Mathematics at Rice University. It is available at <http://www.caam.rice.edu/software/ARPACK>.
- JDBSYM is a C library implementation of the JACOBI-DAVIDSON method optimized for symmetric eigenvalue problems. It solves eigenproblems of the form $Ax = \lambda x$ and $Ax = \lambda Bx$ with or without preconditioning, where A is symmetric and B is symmetric positive definite. It has been developed at the Computer Science Department of the ETH Zurich and is available at <http://people.web.psi.ch/geus/software.html>.
- PARDISO is a fast direct solver package, developed at the Computer Science Department of the University of Basel and available at <http://www.computational.unibas.ch/cs/scicomp/software>.
- ILUPACK is an algebraic multilevel preconditioning software package. This iterative solver has been developed at the Mathematics Department of the Technical University of Berlin. It is available at <http://www.math.tu-berlin.de/ilupack>.

8.2. Numerical Results. In our numerical experiments we will first compare the classical CWI with the shift-and-invert Lanczos method using implicit restarts. The latter is part of ARPACK [42]. For the solution of the symmetric indefinite system $A - \theta I$ we use the most recent version of sparse direct solver PARDISO [57]. This version is based on symmetric weighted matchings and uses METIS as symmetric reordering strategy. The numerical results deal with the computation of 5 eigenvalues of the Anderson matrix A near $\lambda = 0$. Here we state the results for the physically most interesting critical disorder strength $w = w_c = 16.5$ (cf. Table 8.1). As can be seen from Table 8.1, the PARDISO-based shift-and-invert Lanczos is clearly superior to the classic CWI method by at least one order of magnitude regarding computation time. Despite this success, with increasing problem size the amount of memory consumed by the sparse direct solver becomes significant and numerical results N larger than 1'000'000 are skipped.

TABLE 8.1

CPU times in seconds and memory requirements in GB to compute at $w = 16.5$ five eigenvalues closest to $\lambda = 0$ of an Anderson matrix of size $N = M^3 \times M^3$ with CWI and ARPACK-PARDISO. For CWI and $M = 90, 100$, not all 5 eigenvalues converged successfully, so the eigenvector reconstruction finished quicker, leading to apparently shorter CPU times ().*

M	N	CWI		ARPACK-PARDISO	
		time	mem.	time	mem.
30	27'000	21	0.01	9	0.08
40	64'000	300	0.02	46	0.28
50	125'000	1'246	0.04	157	0.68
60	216'000	4'748	0.07	495	1.49
70	343'000	15'100	0.11	1'309	3.00
80	512'000	39'432	0.16	3'619	5.12
90	729'000	97'119*	0.23	7'909	8.70
100	1'000'000	255'842*	0.32	20'239	14.34

Instead, we switch to the ILUPACK-based preconditioner that is also based on symmetric weighted matchings and in addition uses inverse-based pivoting. In particular, for our experiments we use $\kappa = 5$ as bound for the norm $\|L^{-1}\|$ of the inverse triangular factor and AMD for the symmetric reordering. We also tried to use METIS but for this particular matrix problem we find that AMD is clearly more memory efficient. Next we compare PARDISO-based shift-and-invert Lanczos (ARPACK) with that using ILUPACK and the simplified QMR as inner iterative solver. Here we use $\varepsilon = 1/\sqrt{N}$ with aggressive dropping and the QMR method is stopped once the norm of residual satisfies $\|Ax - b\| \leq 10^{-10}\|b\|$. In order to illustrate the benefits of using symmetric weighted matchings we also tried ILUPACK without matching, but the numerical results are disappointing as can be seen from the †s in Table 8.2. We emphasize that the multi-level approach is crucial, a simple use of incomplete factorization methods without multi-level preconditioning [38] does not give the desired results. Besides the effect of matchings we also compare how the performance of the methods changes when varying the value w from the critical value $w = w_c = 16.5$ to $w = 12.0$ and $w = 21.0$. We find that these changes do not affect the sparse direct solver at all while the multilevel ILU significantly varies in its performance. Up to now our explanation for this effect is the observation that with increasing w the diagonal dominance of the system also increases and the ILUPACK preconditioner gains from higher diagonal dominance. As we can see from Table 8.2, ILUPACK still uses

significantly less memory than the direct solver PARDISO for all values of w and it is the only method we were able to use for larger N due to the memory constraints. Also, the computation time is best.

TABLE 8.2

CPU times in seconds and memory requirements in GB to compute five eigenvalues closest to $\lambda = 0$ of an Anderson matrix of size $N^3 \times N^3$ with ARPACK-PARDISO, ARPACK-ILUPACK, and ARPACK-ILUPACK-SYMMATCH. The symbol ‘—’ indicates that a memory consumption was larger than 25 GB and ‘†’ indicates memory problems with respect to the fill-in.

M	W	ARPACK					
		PARDISO		ILUPACK		ILUPACK-SYMMATCH	
		time	mem.	time	mem.	time	mem.
70	12.0	1’359	3.00	5’117	1.09	2’140	0.95
100	12.0	20’639	14.34	39’222	5.62	13’583	3.20
130	12.0	—	—	†	†	65’722	
70	16.5	1’305	3.00	504	0.33	477	0.31
100	16.5	20’439	14.34	2’349	0.95	2’177	0.89
130	16.5	—	—	6’320	2.09	6’530	1.95
160	16.5	—	—	23’663	3.95	13’863	3.63
70	21.0	1’225	3.00	371	0.22	310	0.22
100	21.0	20’239	14.34	1’513	0.64	1’660	0.65
130	21.0	—	—	3’725	1.41	3’527	1.44
160	21.0	—	—	15’302	2.63	20’120	2.68

When using preconditioning methods inside shift-and-invert Lanczos we usually have to solve the inner linear system for $A - \theta I$ up to the machine precision to make sure that the eigenvalues and eigenvectors are sufficiently correct. In contrast to this the JACOBI-DAVIDSON method allows to solve the associated correction equation less accurately and only when convergence takes place a more accurate solution is required. In order to show the significant difference between the iterative parts of ARPACK and JACOBI-DAVIDSON we state the number of iteration steps in Table 8.3. If we were to aim for more eigenpairs, we expect that eventually the JDBSYM becomes less efficient and should again be replaced by ARPACK.

In the sequel we compare the traditional CWI method with the JACOBI-DAVIDSON code JDBSYM [33] using ILUPACK as preconditioner. Table 8.4 shows that switching from ARPACK to JACOBI-DAVIDSON in this case improves the total method by another factor 6 or greater. For this reason JACOBI-DAVIDSON together with ILUPACK will be used as default solver in the following. The numerical results in Table 8.4 show a dramatic improvement in the computation time by using ILUPACK-based JACOBI-DAVIDSON. Although this new method slows down for smaller w due to poorer diagonal dominance, a gain by orders of magnitude can still be observed. For $w = 16.5$ and larger, even more than three orders of magnitude in the computation time can be observed. Hence the new method drastically outperforms the CWI method while the memory requirement is still moderate.

One key to the success of the preconditioner is based on the threshold κ which bounds the growth of L^{-1} . Already for a small example such as $M = 70$ significant differences can be observed. As we show in Table 8.5, increasing the bound by a factor 2 from $\kappa = 5$ up to $\kappa = 10$ and $\kappa = 20$ leads to an enormous increase of fill. Here we measure the fill of the incomplete LDL^T factorization relative to the non-zeros of the original matrix. By varying the drop tolerance ε we also see that the dependence

TABLE 8.3

Number of inner/outer interaction steps inside ARPACK and JACOBI-DAVIDSON. The symbol '—' indicates that the computations were not performed anymore for ARPACK.

M	W	ILUPACK-SYMMATCH					
		ARPACK			JACOBI-DAVIDSON		
		outer	total	inner average	outer	total	inner average
70	12.0	42	871	20.7	43	246	5.7
100	12.0	43	1101	25.6	44	325	7.4
130	12.0	42	1056	25.1	44	274	6.2
70	16.5	43	611	14.2	41	200	4.88
100	16.5	43	857	19.9	43	231	5.37
130	16.5	42	1058	25.2	38	313	8.24
160	16.5	42	968	23.1	41	276	6.73
190	16.5	—	—	—	39	339	8.69
220	16.5	—	—	—	40	433	10.82
250	16.5	—	—	—	47	652	13.87
70	21.0	43	585	13.60	40	200	5.00
100	21.0	42	1004	23.90	42	301	7.17
130	21.0	44	914	20.77	39	274	7.03
160	21.0	25	896	35.84	43	507	11.79
190	21.0	—	—	—	46	637	13.85
220	21.0	—	—	—	41	855	20.85
250	21.0	—	—	—	43	891	20.72

on κ is much more significant than the dependence of ε . Roughly speaking, the ILU decomposition becomes twice as expensive when κ is replaced by 2κ and so does the fill-in. The latter is crucial since memory constraints severely limit the size of the application that can be computed.

In Table 8.6 we show how JDBSYM and ILUPACK-SYMMATCH perform when instead of periodic boundary conditions, we use hard wall boundaries, i.e., $x_{0;j;k} = x_{i;0;k} = x_{i;j;0} = x_{M+1;j;k} = x_{i;M+1;k} = x_{i;j;M+1} = 0$ for all i, j, k . This is sometimes of interest in the Anderson problem and generally, it is expected that for large M , the difference in eigenvalues and eigenvectors becomes small when compared to the standard periodic boundaries. In addition, we also show results for the so-called off-diagonal Anderson problem [14]. Here, we shift the diagonal to a constant $\sigma = 1.28$ and incorporate the randomness by setting the off-diagonal elements of A to be uniformly distributed in $[-1/2, 1/2]$. The graph of the matrix A remains the same. These values correspond — similarly to $w_c = 16.5$ used before for purely diagonal randomness — to the physically most interesting localization transition in this model [14]. We note that using hard wall boundary conditions instead of periodic boundary conditions leads to slightly less fill but increases the number of iteration steps as can be seen in Table 8.6. This conclusion carries over to the off-diagonal Anderson problem, where the memory consumption is less but the iterative part takes even longer. In principle our results could be improved if we were to switch to a smaller threshold ε than the here uniformly applied $\varepsilon = 1/\text{sqrt}(N)$.

TABLE 8.4

CPU times in seconds and memory requirements in GB to compute five eigenvalues closest to $\lambda = 0$ with CWI and JACOBI-DAVIDSON using ILUPACK-SYMMATCH for the shift-and-invert technique. ‡' indicates that the convergence of the method was too slow. For CWI and $M = 100$, not all 5 eigenvalues converged successfully, so the eigenvector reconstruction finished quicker, leading to variances in the CPU times ().*

M	W	CWI		JACOBI-DAVIDSON ILUPACK-SYMMATCH	
		time	mem.	time	mem.
70	12.0	20'228	0.11	1'314	0.95
100	12.0	148'843	0.32	8'522	2.93
130	12.0	‡	‡	56'864	8.06
70	16.5	15'100	0.11	258	0.34
100	16.5	255'842*	0.32	978	0.98
130	16.5	‡	‡	2'895	2.16
160	16.5	‡	‡	5'860	4.05
190	16.5	‡	‡	16'096	6.75
220	16.5	‡	‡	30'160	10.58
250	16.5	‡	‡	86'573	
70	21.0	14'371	0.11	255	0.26
100	21.0	331'514*	0.32	1'134	0.75
130	21.0	‡	‡	1'565	1.65
160	21.0	‡	‡	4'878	3.08
190	21.0	‡	‡	11'032	5.22
220	21.0	‡	‡	28'334	
250	21.0	‡	‡	48'223	

TABLE 8.5

The influence of the inverse bound κ on the amount of memory. For $M = 70$ compare for different thresholds how the fill-in $\text{nnz}(LDL^T)/\text{nnz}(A)$ varies depending on κ and state the computation time in seconds.

ε	$\kappa = 5$			$\kappa = 10$			$\kappa = 20$		
	fill	time LDL^T	total time	fill	time LDL^T	total time	fill	time LDL^T	total time
0.01	5.4	$3.7 \cdot 10^1$	$8.7 \cdot 10^2$	8.7	$6.7 \cdot 10^1$	$5.0 \cdot 10^2$	15.2	$1.6 \cdot 10^2$	$4.8 \cdot 10^2$
0.005	6.8	$5.4 \cdot 10^1$	$4.4 \cdot 10^2$	11.0	$1.0 \cdot 10^2$	$3.8 \cdot 10^2$	19.1	$2.3 \cdot 10^2$	$5.0 \cdot 10^2$
0.0025	8.6	$8.1 \cdot 10^1$	$3.1 \cdot 10^2$	13.8	$1.5 \cdot 10^2$	$3.6 \cdot 10^2$	24.1	$3.4 \cdot 10^2$	$6.0 \cdot 10^2$
0.001	11.7	$1.3 \cdot 10^2$	$3.0 \cdot 10^2$	18.0	$2.3 \cdot 10^2$	$4.1 \cdot 10^2$	32.1	$5.4 \cdot 10^2$	$7.8 \cdot 10^2$

9. Conclusion. We have shown that modern approaches to preconditioning based on symmetric matchings and multilevel preconditioning methods lead to an astonishing increase in performance and available system sizes for the Anderson model of localization. This approach is not only several orders of magnitudes faster than the traditional CWI approach, it also consumes only a moderate amount of memory thus allowing to study the Anderson eigenproblem for significantly larger scales than ever before.

Let us briefly recall the main ingredients necessary for this progress: At the heart of the new approach lies the use of symmetric matchings [38] in the preconditioning

TABLE 8.6

Difference in performance for our standard problem with periodic boundary conditions, the problem with hard wall conditions, and the inverse problem with random numerical entries in the off-diagonal elements. Memory requirement (in GB) and CPU times (in seconds) to compute at the transition the eigenvectors corresponding to the five eigenvalues closest to $\lambda = 0$ with shift-and-invert JACOBI-DAVIDSON and the ILUPACK-SYMMATCH solver using symmetric weighted matchings.

N	Periodic		Hard wall		Inverse	
	time	memory	time	memory	time	memory
70	258	0.34	282	0.31	457	0.27
100	980	0.98	969	0.94	2'075	0.79
130	5'244	2.16	2'090	2.07	6'472	1.76
160	9'958	4.05	5'661	3.90	11'975	3.30
190	14'742	6.75	13'431	6.62	27'488	

stage of the inverse-based incomplete factorization preconditioning iterative method [9]. Furthermore, the preconditioning itself is of a multi-level type, complementary to the often used full-pivoting strategies. Next, the inverse-based approach is also of paramount importance to keep the fill-in at a manageable level (see Table 8.5). And last, we emphasize that these results, of course, reflect our selected problem class: to compute a few of the interior eigenvalues and associated eigenvectors for a highly indefinite symmetric matrix defined by the Anderson model of localization.

The performance increase by several orders of magnitude (see Table 8.4) is solely due to our use of new and improved algorithms. Combined with advances in the performance to cost ratio of computing hardware during the last 6 years period, current preconditioning methods makes it possible to solve those problems quickly and easily which have been considered by far too large until recently [13]. Even for $N \times N$ matrices as large as $N = 64 \cdot 10^6$, it is now possible within a few days to compute the interior eigenstates of the Anderson problem.

The success of this method indicates that it might also be successfully applied to other large-scale problems arising in (quantum) physics.

Acknowledgement. We gratefully acknowledge discussions and implementation help from A. Croy and C. Sohrmann in the initial stages of the project.

REFERENCES

- [1] P. AMESTOY, T. A. DAVIS, AND I. S. DUFF, *An approximate minimum degree ordering algorithm*, SIAM J. Matrix Analysis and Applications, 17 (1996), pp. 886–905.
- [2] P. W. ANDERSON, *Absence of diffusion in certain random lattices*, Phys. Rev., 109 (1958), pp. 1492–1505.
- [3] H. AOKI, *Fractal dimensionality of wave functions at the mobility edge: Quantum fractal in the Landau levels*, Phys. Rev. B, 33 (1986), pp. 7310–7313.
- [4] P. ARBENZ AND R. GEUS, *Multilevel preconditioned iterative eigensolvers for maxwell eigenvalue problems*, Appl. Numer. Math., 54 (2005), pp. 107–121.
- [5] M. BENZI, *Preconditioning techniques for large linear systems: A survey*, Journal of Computational Physics, 182 (2002), pp. 418–477.
- [6] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numerica, 14 (2005), pp. 1–137.
- [7] M. BENZI, J. HAWS, AND M. TÛMA, *Preconditioning highly indefinite and nonsymmetric matrices*, SIAM J. Scientific Computing, 22 (2000), pp. 1333–1353.
- [8] M. BOLLHÖFER AND Y. SAAD, *Multilevel preconditioners constructed from inverse-based ILUs*, SIAM J. Scientific Computing, to appear (2005).

- [9] M. BOLLHÖFER AND O. SCHENK, *ILUPACK volume 2.0 — preconditioning software package for symmetrically structured problems*, May 2005. <http://www.math.tu-berlin.de/ilupack/>.
- [10] T. BRANDES, B. HUCKESTEIN, AND L. SCHWEITZER, *Critical dynamics and multifractal exponents at the Anderson transition in 3d disordered systems*, *Ann. Phys. (Leipzig)*, 5 (1996), pp. 633–651.
- [11] J. R. BUNCH, *Partial pivoting strategies for symmetric matrices*, *SIAM J. Numerical Analysis*, 11 (1974), pp. 521–528.
- [12] J. R. BUNCH AND L. KAUFMAN, *Some stable methods for calculating inertia and solving symmetric linear systems*, *Mathematics of Computation*, 31 (1977), pp. 163–179.
- [13] P. CAIN, F. MILDE, R. A. RÖMER, AND M. SCHREIBER, *Use of cluster computing for the anderson model of localization*, *Comp. Phys. Comm.*, 147 (2002), pp. 246–250.
- [14] P. CAIN, R. A. RÖMER, AND M. SCHREIBER, *Phase diagram of the three-dimensional Anderson model of localization with random hopping*, *Ann. Phys. (Leipzig)*, 8 (1999), pp. SI33–SI38. ArXiv: cond-mat/9908255.
- [15] A. CLINE, C. B. MOLER, G. STEWART, AND J. WILKINSON, *An estimate for the condition number of a matrix*, *SIAM J. Numerical Analysis*, 16 (1979), pp. 368–375.
- [16] J. CULLUM AND R. A. WILLOUGHBY, *Lanczos Algorithms for Large Symmetric Eigenvalue Computations, Volume 1: Theory*, Birkhäuser, Boston, 1985.
- [17] ———, *Lanczos Algorithms for Large Symmetric Eigenvalue Computations, Volume 2: Programs*, Birkhäuser, Boston, 1985. <http://www.netlib.org/lanczos/>.
- [18] P. DAYAL, M. TROYER, AND R. VILLIGER, *The Iterative Eigensolver Template Library*, ETH Zürich. <http://www.comp-phys.org:16080/software/ietl/>.
- [19] J. W. DEMMEL, S. C. EISENSTAT, J. R. GILBERT, X. S. LI, AND J. W. H. LIU, *A supernodal approach to sparse partial pivoting*, *SIAM J. Matrix Analysis and Applications*, 20 (1999), pp. 720–755.
- [20] D. DODSON AND J. G. LEWIS, *Issues relating to extension of the basic linear algebra subprograms*, *ACM SIGNUM Newslett.*, 20 (1985), pp. 2–18.
- [21] J. J. DONGARRA, J. DU CROZ, S. HAMMARLING, AND R. J. HANSON, *Issues relating to extension of the basic linear algebra subprograms*, *ACM SIGNUM Newslett.*, 20 (1985), pp. 2–18.
- [22] I. S. DUFF AND J. R. GILBERT, *Symmetric weighted matching for indefinite systems*. Householder Symposium XV, 2002, Peebles, June 2002. talk.
- [23] I. S. DUFF AND J. KOSTER, *The design and use of algorithms for permuting large entries to the diagonal of sparse matrices*, *SIAM J. Matrix Analysis and Applications*, 20 (1999), pp. 889–901.
- [24] I. S. DUFF AND S. PRALET, *Strategies for scaling and pivoting for sparse symmetric indefinite problems*, Technical Report TR/PA/04/59, CERFACS, 2004.
- [25] I. S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear equations*, *ACM Trans. Math. Softw.*, 9 (1983), pp. 302–325.
- [26] A. EILMES, R. A. RÖMER, AND M. SCHREIBER, *The two-dimensional Anderson model of localization with random hopping*, *Eur. Phys. J. B*, 1 (1998), pp. 29–38.
- [27] U. ELSNER, V. MEHRMANN, F. MILDE, R. A. RÖMER, AND M. SCHREIBER, *The Anderson model of localization: a challenge for modern eigenvalue methods*, *SIAM J. Sci. Comp.*, 20 (1999), pp. 2089–2102. ArXiv: physics/9802009.
- [28] R. D. FOKKEMA, G. L. G. SLEIJPEN, AND H. A. V. DER VORST, *Jacobi–Davidson style QR and QZ algorithms for the reduction of matrix pencils*, *SIAM J. Scientific Computing*, 20 (2000), pp. 94–125.
- [29] R. FREUND AND F. JARRE, *A QMR–based interior–point algorithm for solving linear programs*, *Mathematical Programming, Series B*, 76 (1997), pp. 183–210.
- [30] R. FREUND AND N. NACHTIGAL, *Software for simplified lanczos and QMR algorithms*, *Appl. Numer. Math.*, 19 (1995), pp. 319–341.
- [31] J. A. GEORGE AND E. NG, *An implementation of Gaussian elimination with partial pivoting for sparse systems*, *SIAM J. Scientific and Statistical Computing*, 6 (1985), pp. 390–409.
- [32] R. GEUS, *Jdbsym version 0.14*. <http://www.inf.ethz.ch/personal/geus/software.html>.
- [33] ———, *The Jacobi–Davidson algorithm for solving large sparse symmetric eigenvalue problems with application to the design of accelerator cavities*, PhD thesis, ETH Zürich, Department Informatik, Institut f. Wissenschaftliches Rechnen, 2002.
- [34] J. R. GILBERT AND E. NG, *Predicting structure in nonsymmetric sparse matrix factorizations*, in *Graph Theory and Sparse Matrix Computation*, J. A. George, J. R. Gilbert, and J. W. H. Liu, eds., Springer-Verlag, 1993.
- [35] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, third ed., 1996.
- [36] U. GRIMM, R. A. RÖMER, AND G. SCHLIECKER, *Electronic states in topologically disordered*

- systems*, Ann. Phys. (Leipzig), 7 (1998), pp. 389–393.
- [37] A. GUPTA AND L. YING, *A fast maximum-weight-bipartite-matching algorithm for reducing pivoting in sparse gaussian elimination*, Tech. Report RC 21576 (97320), IBM T. J. Watson Research Center, Yorktown Heights, NY, October 1999.
- [38] M. HAGEMANN AND O. SCHENK, *Weighted matchings for the preconditioning of symmetric indefinite linear systems*, Technical Report CS-2004-005, Department of Computer Science, University of Basel, 2004.
- [39] G. KARYPIS AND V. KUMAR, *A fast and high quality multilevel scheme for partitioning irregular graphs*, SIAM J. Scientific Computing, 20 (1998), pp. 359–392.
- [40] B. KRAMER, A. BRODERIX, A. MACKINNON, AND M. SCHREIBER, *The Anderson transition: new numerical results for the critical exponents*, Physica A, 167 (1990), pp. 163–174.
- [41] B. KRAMER AND A. MACKINNON, *Localization: theory and experiment*, Rep. Prog. Phys., 56 (1993), pp. 1469–1564.
- [42] R. B. LEHOUCQ, K. MASHHOFF, D. C. SORENSEN, AND C. YANG, *Arpack*. <http://www.caam.rice.edu/software/ARPACK/>.
- [43] R. B. LEHOUCQ, D. C. SORENSEN, AND C. YANG, *ARPACK users' guide*, SIAM Publications, Philadelphia, 1998.
- [44] N. LI, Y. SAAD, AND E. CHOW, *Crout versions of ILU for general sparse matrices*, SIAM J. Scientific Computing, 25 (2004), pp. 716–728.
- [45] Q. LI, S. KATSOPRINAKIS, E. N. ECONOMOU, AND C. M. SOUKOULIS, *Scaling properties in highly anisotropic systems*, Phys. Rev. B, 56 (1997), pp. R4297–R4300. ArXiv: cond-mat/9704104.
- [46] X. S. LI AND J. W. DEMMEL, *SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems*, ACM Trans. Math. Softw. Accepted, in press.
- [47] F. MILDE, R. A. RÖMER, AND M. SCHREIBER, *Multifractal analysis of the metal-insulator transition in anisotropic systems*, Phys. Rev. B, 55 (1997), pp. 9463–9469.
- [48] ———, *Energy-level statistics at the metal-insulator transition in anisotropic systems*, Phys. Rev. B, 61 (2000), pp. 6028–6035.
- [49] M. OLSCHOWKA AND A. NEUMAIER, *A new pivoting strategy for gaussian elimination*, Linear Algebra and its Applications, 240 (1996), pp. 131–151.
- [50] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numerical Analysis, 12 (1975), pp. 617–629.
- [51] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, Prentice-Hall, 1980.
- [52] I. PLYUSHCHAY, R. A. RÖMER, AND M. SCHREIBER, *The three-dimensional anderson model of localization with binary random potential*, Phys. Rev. B, 68 (2003), pp. 064201–8.
- [53] D. PORATH, G. CUNIBERTI, AND R. DI FELICE, *Charge transport in DNA-based devices*, Topics in Current Chemistry, 237 (2004), p. 183.
- [54] R. A. RÖMER AND M. SCHREIBER, *The Anderson Transition and its Ramifications — Localisation, Quantum Interference, and Interactions*, Springer, Berlin, 2003, ch. Numerical investigations of scaling at the Anderson transition, pp. 3–19.
- [55] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM Publications, Philadelphia, second ed., 2003.
- [56] Y. SAAD AND M. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Scientific and Statistical Computing, 7 (1986), pp. 856–869.
- [57] O. SCHENK AND K. GÄRTNER, *On fast factorization pivoting methods for symmetric indefinite systems*, Technical Report, Computer Science Department, University of Basel, Switzerland, 2004. Submitted to Electronic Transactions on Numerical Analysis.
- [58] ———, *Solving unsymmetric sparse systems of linear equations with PARDISO*, Journal of Future Generation Computer Systems, 20 (2004), pp. 475–487.
- [59] O. SCHENK, S. RÖLLIN, AND A. GUPTA, *The effects of unsymmetric matrix permutations and scalings in semiconductor device and circuit simulation*, IEEE Transactions On Computer-Aided Design Of Integrated Circuits And Systems, 23 (2004).
- [60] M. SCHREIBER AND M. OTTOMEIER, *Localization of electronic states in 2D disordered systems*, J. Phys.: Condens. Matter, 4 (1992), pp. 1959–1971.
- [61] G. L. G. SLEIJPEN AND H. A. V. DER VORST, *A Jacobi–Davidson iteration for linear eigenvalue problems*, SIAM J. Matrix Analysis and Applications, 17 (1996), pp. 401–425.
- [62] D. SORENSEN, *Implicit application of polynomial filters in a k-step Arnoldi method*, SIAM J. Matrix Analysis and Applications, 13 (1992), pp. 357–385.
- [63] C. M. SOUKOULIS AND E. N. ECONOMOU, *Off-diagonal disorder in one-dimensional systems*, Phys. Rev. B, 24 (1981), p. 5698.
- [64] ———, *Fractal character of eigenstates in disordered systems*, Phys. Rev. Lett., 52 (1984),

