

8 The Lanczos Method

Erik Koch

Computational Materials Science

German Research School for Simulation Sciences

Contents

1	Lanczos Method	3
1.1	Krylov space	5
1.2	Spectral functions	8
2	Application to the Hubbard model	11
2.1	Representation of basis and Hamiltonian	13
2.2	Green functions	14
2.3	Parallelization strategies	16
3	Application to DMFT	18
3.1	Cluster methods	21
3.2	Anderson impurity model	21
3.3	Hybridization sum-rules	23
3.4	Symmetries	26
4	Conclusions	29

The Lanczos iteration [1] was conceived as a method for tridiagonalizing Hermitian matrices. Like the related Arnoldi method [2] for non-Hermitian matrices, it initially received widespread attention. Iterative approaches were, however, soon eclipsed by direct methods (Householder transformations and QR factorization), that are better suited for solving the eigenvalue problem for general matrices. Actually, the Lanczos method is particularly suited for the determination of extreme eigenvalues and -vectors. Therefore, it was rediscovered in the 1970s [3], when computers had become sufficiently powerful to treat matrices large enough for the Lanczos algorithm to outperform general methods, nicely illustrating the Fundamental Law of Computer Science: the faster the computer, the greater the importance of the speed of algorithms [4]. By now iterative methods are an integral part of the numerical linear algebra curriculum [4–6].

For finding eigenvalues of a matrix H of dimension N , the Lanczos method requires the evaluation of matrix-vector products $H \cdot v$ as the only problem-specific step. This matrix-vector product can be calculated particularly efficiently when the matrix H is sparse, i.e., when the number of non-zero matrix elements per row does not scale with the matrix dimension. Storing such a matrix takes only $\mathcal{O}(N)$ memory and $H \cdot v$ can be evaluated in $\mathcal{O}(N)$ time. Calculating the extremal eigenvalues requires $\mathcal{O}(1)$ iterations, i.e., overall $\mathcal{O}(N)$ time. For comparison, a direct diagonalization takes $\mathcal{O}(N^2)$ for storing the matrix and $\mathcal{O}(N^3)$ time to diagonalize. Besides their favorable scaling for sparse matrix problems, iterative methods have the advantage that they systematically approach the desired result. Typically the iteration converges geometrically and can be stopped as soon as the desired accuracy is reached. In contrast, direct methods appear to make no progress towards the solution until all $\mathcal{O}(N^3)$ operations are completed and the full result is obtained.

Since the Lanczos method is particularly suited for dealing with large sparse Hamiltonians, it is the method of choice for systems with short-range interactions. For band-structure calculations in a linear combination of atomic orbitals (LCAO) or tight-binding (TB) basis, it is known as the recursion method [7]. The basic idea here is to switch from the Bloch picture of a perfectly periodic solid to a local picture, replacing the solution of the Schrödinger equation in terms of Bloch waves by the calculation of the local density of states. The crucial technical point is to calculate the density of states not via a spectral representation (in terms of Bloch waves), but by repeated application of the Hamiltonian H to a localized single-electron state. With each application of H the electron explores more and more sites. Thus, if the hopping matrix elements beyond a certain distance are zero, such calculations can be performed without having to restrict the system to finite size.

For many-body models like quantum-spin- or Hubbard-models [8] this is unfortunately not possible. They have to be defined on a finite cluster, giving rise to a finite-dimensional Hamiltonian matrix. Since the size of the Hilbert space grows exponentially with system-size, actual calculations are restricted by the available computer memory. In a typical simulation, first the ground-state is calculated by a Lanczos iteration. Building on this, spectral functions are calculated in a similar way as in the recursion method. The great advantage of this approach is that it gives the dynamical properties of the ground state ($T=0$) directly on the real axis. The price is the restriction to (small) finite-size systems.

1 Lanczos Method

We can find the ground-state $|\Psi_0\rangle$ and its energy E_0 for a Hamiltonian H from the variational principle. The wavefunction-functional

$$E[\Psi] = \frac{\langle \Psi | H | \Psi \rangle}{\langle \Psi | \Psi \rangle} \quad (1)$$

is minimized for $\Psi = \Psi_0$, with $E[\Psi_0] = E_0$. The functional gradient

$$\frac{\delta E[\Psi]}{\delta \langle \Psi |} = \frac{H|\Psi\rangle - E[\Psi]|\Psi\rangle}{\langle \Psi | \Psi \rangle} = |\Psi_a\rangle \quad (2)$$

gives the direction of steepest-ascent of the functional from the point $|\Psi\rangle$. Moving in the opposite direction will thus result in a wavefunction with lower energy expectation value: $E[\Psi - \alpha\Psi_a] < E[\Psi]$ for small, positive α .

To find the optimum value of α , we minimize $E[\Psi - \alpha\Psi_a]$. For this, it is convenient to introduce an orthogonal basis in the space spanned by the two vectors $|\Psi\rangle$ and $|\Psi_a\rangle$. From (2) we see that $\text{span}(|\Psi\rangle, |\Psi_a\rangle) = \text{span}(|\Psi\rangle, H|\Psi\rangle)$. As first basis vector, we normalize $|\Psi\rangle$

$$|v_0\rangle = |\Psi\rangle / \sqrt{\langle \Psi | \Psi \rangle},$$

for the second vector we orthogonalize $H|v_0\rangle$ to $|v_0\rangle$

$$|\tilde{v}_1\rangle = H|v_0\rangle - |v_0\rangle\langle v_0 | H | v_0 \rangle \quad (3)$$

and normalize to obtain $|v_1\rangle$. With $a_n = \langle v_n | H | v_n \rangle$ and $b_1^2 = \langle \tilde{v}_1 | \tilde{v}_1 \rangle$ we thus have

$$H|v_0\rangle = b_1|v_1\rangle + a_0|v_0\rangle \quad (4)$$

from which we see that $\langle v_1 | H | v_0 \rangle = b_1$.

We can then write any normalized wavefunction in $\text{span}(|\Psi\rangle, H|\Psi\rangle) = \text{span}(|v_0\rangle, |v_1\rangle)$ as

$$|v\rangle = \cos(\theta)|v_0\rangle + \sin(\theta)|v_1\rangle. \quad (5)$$

Minimizing the expectation value

$$\langle v | H | v \rangle = a_0 \cos^2(\theta) + 2b_1 \sin(\theta) \cos(\theta) + a_1 \sin^2(\theta), \quad (6)$$

with respect to θ , we obtain, dividing by $\cos^2(\theta)$, the quadratic equation

$$b_1 \tan^2(\theta) + (a_0 - a_1) \tan(\theta) - b_1 = 0. \quad (7)$$

Solving for θ we find the lowest-energy state on the subspace spanned by $|v_0\rangle$ and $H|v_0\rangle$. Alternatively, we can diagonalize the Hamiltonian matrix on the two-dimensional subspace, which in the basis $|v_0\rangle, |v_1\rangle$ is given by

$$H_{\text{span}(|\Psi\rangle, H|\Psi\rangle)} = \begin{pmatrix} a_0 & b_1 \\ b_1 & a_1 \end{pmatrix}. \quad (8)$$

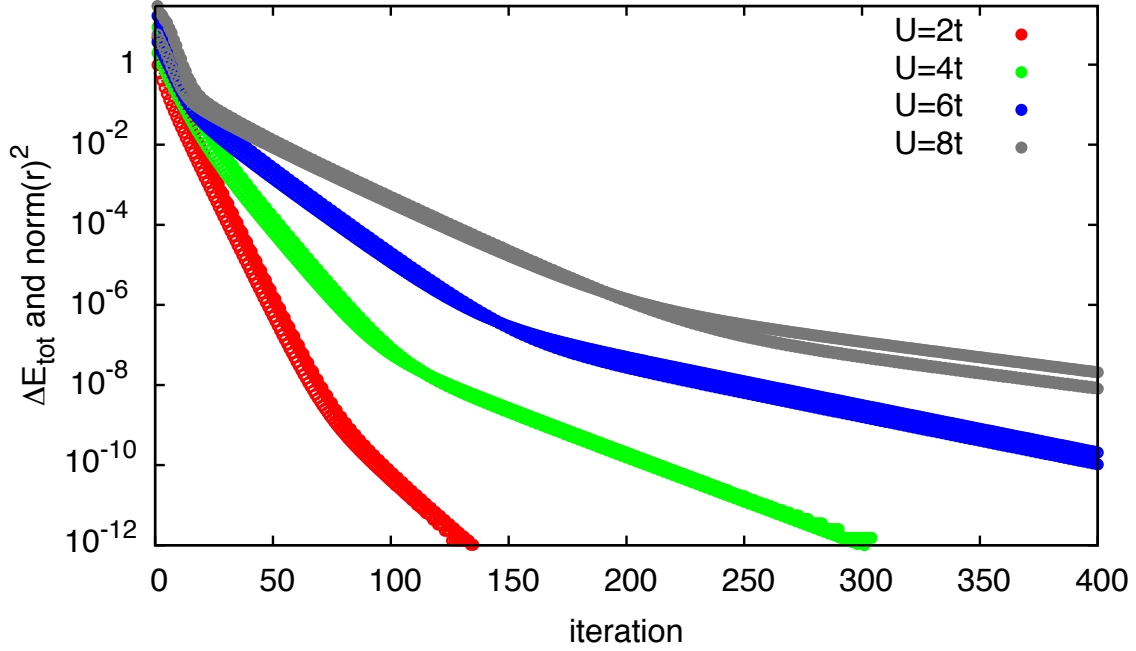


Fig. 1: Convergence of the residual (filled circles) and the lowest eigenvalue (open circles) for a steepest-descent minimization of a Hubbard-chain of 10 sites at half-filling, starting from a random initial vector.

Naturally, we can use the variational state of lowest energy

$$|\Psi^{(2)}\rangle = \cos(\theta_{\min})|v_0\rangle + \sin(\theta_{\min})|v_1\rangle \quad (9)$$

as the starting point for another steepest-descent minimization. Doing this repeatedly, we obtain a series of vectors with decreasing energy expectation value, which rapidly converge to a minimum. For a generic functional, this would usually be a local, not the global minimum, which makes the optimization of high-dimensional functions a hard problem. The energy functional (1), however, has only minima for the ground-states, all other stationary points are saddle points. We can thus expect rapid convergence to the ground state, examples given in figure 1, except in the case where the the gradient (2) vanishes, i.e., if $|\Psi\rangle$ happens to be an eigenfunction of H .

For checking convergence of this steepest-descent method, introduced by Kantorovich [9] and, independently, by Hestenes and Karush [10], we can monitor the change in the energy expectation value or determine when the residual

$$r[\Psi] = \|(H - E[\Psi])|\Psi\rangle\|^2 = \langle\Psi|H^2|\Psi\rangle - E[\Psi]^2, \quad (10)$$

which measures the quality of the eigenstate, becomes sufficiently small.

1.1 Krylov space

If we apply the method of steepest-descent L times, starting from a vector $|v_0\rangle$, the resulting vector will lie in $\mathcal{K}^L(|v_0\rangle) = \text{span}(|v_0\rangle, H|v_0\rangle, H^2|v_0\rangle, \dots, H^L|v_0\rangle)$, the $L + 1$ -dimensional *Krylov space* [11] of H over $|v_0\rangle$. Instead of repeatedly minimizing the energy in two-dimensional subspaces, we could directly find the state of lowest energy in $\mathcal{K}^L(|v_0\rangle)$. Having more degrees of freedom for the minimization will lead to even faster convergence.

To implement this idea, we construct an orthonormal basis $|v_n\rangle$ of the Krylov space. We start with the normalized vector $|v_0\rangle$. The second basis vector $|v_1\rangle$ is constructed as in the steepest-descent method (3):

$$b_1|v_1\rangle = |\tilde{v}_1\rangle = H|v_0\rangle - a_0|v_0\rangle. \quad (11)$$

The next basis vector is likewise constructed as $H|v_n\rangle$ orthogonalized to all previous vectors, and normalized

$$b_2|v_2\rangle = |\tilde{v}_2\rangle = H|v_1\rangle - \sum_{i=0}^1 |v_i\rangle\langle v_i|H|v_1\rangle = H|v_1\rangle - a_1|v_1\rangle - b_1|v_0\rangle. \quad (12)$$

where $a_n = \langle v_n|H|v_n\rangle$ and $b_n^2 = \langle \tilde{v}_n|\tilde{v}_n\rangle$. The fourth basis vector is

$$b_3|v_3\rangle = |\tilde{v}_3\rangle = H|v_2\rangle - \sum_{i=0}^2 |v_i\rangle\langle v_i|H|v_2\rangle = H|v_2\rangle - a_2|v_2\rangle - b_2|v_1\rangle. \quad (13)$$

Here the last term in the orthogonalization vanishes, because (11) together with the orthogonality of the basis vectors for $n = 0 \dots 2$ implies $\langle v_2|H|v_0\rangle = 0$.

The construction of the further basis vectors follows the same scheme

$$b_{n+1}|v_{n+1}\rangle = |\tilde{v}_{n+1}\rangle = H|v_n\rangle - \sum_{i=0}^n |v_i\rangle\langle v_i|H|v_n\rangle = H|v_n\rangle - a_n|v_n\rangle - b_n|v_{n-1}\rangle$$

with $a_n = \langle v_n|H|v_n\rangle$ and $b_n^2 = \langle \tilde{v}_n|\tilde{v}_n\rangle$. Rearranging shows that H is tridiagonalized

$$H|v_n\rangle = b_n|v_{n-1}\rangle + a_n|v_n\rangle + b_{n+1}|v_{n+1}\rangle$$

which in turn implies that $H|v_i\rangle$ is orthogonal to all basis states, except $|v_i\rangle$ and $|v_{i\pm 1}\rangle$. This tridiagonalization of H is the essence of the *Lanczos method* [1].

After L steps the Hamiltonian on the $L + 1$ -dimensional Krylov space is given by

$$H_{\mathcal{K}^L(|v_0\rangle)} = \begin{pmatrix} a_0 & b_1 & 0 & 0 & \dots & 0 & 0 \\ b_1 & a_1 & b_2 & 0 & \dots & 0 & 0 \\ 0 & b_2 & a_2 & b_3 & & 0 & 0 \\ 0 & 0 & b_3 & a_3 & & 0 & 0 \\ & \vdots & & \ddots & \ddots & & \\ 0 & 0 & 0 & 0 & & a_{L-1} & b_L \\ 0 & 0 & 0 & 0 & \dots & b_L & a_L \end{pmatrix} \quad (14)$$

v=init	
b0=norm2(v)	not part of tridiagonal matrix
scal(1/b0,v)	v= v ₀ ⟩
w=0	
w=w+H*v	w=H v ₀ ⟩
a[0]=dot(v,w)	
axpy(-a[0],v,w)	w= ṽ ₁ ⟩ = H v ₀ ⟩ - a ₀ v ₀ ⟩
b[1]=norm2(w)	
for n=1,2,...	
if abs(b[n])<eps then exit	invariant subspace
scal(1/b[n],w)	w= v _n ⟩
scal(-b[n],v)	v=-b _n v _{n-1} ⟩
swap(v,w)	
w=w+H*v	w=H v _n ⟩ - b _n v _{n-1} ⟩
a[n]=dot(v,w)	a[n]=⟨v _n H v _n ⟩ - b _n ⟨v _n v _{n-1} ⟩
axpy(-a[n],v,w)	w= ṽ _{n+1} ⟩
b[n+1]=norm2(w)	
diag(a[0]..a[n], b[1]..b[n])	getting a _{n+1} needs another H v⟩
if converged then exit	
end	

Table 1: *The implementation of the Lanczos iteration requires only two N -dimensional vectors for tridiagonalizing H and thus for calculating the ground-state energy. Constructing the Lanczos-approximation of the ground-state vector requires a second iteration and one additional N -dimensional vector. The by far most expensive operation is the matrix-vector product.*

If we do not normalize the basis vectors, we obtain an iteration of the form

$$|\Phi_{n+1}\rangle = H|\Phi_n\rangle - \frac{\langle\Phi_n|H|\Phi_n\rangle}{\langle\Phi_n|\Phi_n\rangle}|\Phi_n\rangle - \frac{\langle\Phi_n|\Phi_n\rangle}{\langle\Phi_{n-1}|\Phi_{n-1}\rangle}|\Phi_{n-1}\rangle \quad (15)$$

where $|\Phi_n\rangle = \prod_{i=1}^n b_i |v_n\rangle$ in terms of which we have

$$a_n = \frac{\langle\Phi_n|H|\Phi_n\rangle}{\langle\Phi_n|\Phi_n\rangle}, \quad b_n^2 = \frac{\langle\Phi_n|\Phi_n\rangle}{\langle\Phi_{n-1}|\Phi_{n-1}\rangle}. \quad (16)$$

In this unnormalized basis the Hamiltonian appears non-Hermitian

$$H|\Phi_n\rangle = b_n^2|\Phi_{n-1}\rangle + a_n|\Phi_n\rangle + |\Phi_{n+1}\rangle, \quad (17)$$

but it actually is

$$\langle\Phi_{n+1}|H|\Phi_n\rangle = \langle\Phi_{n+1}|\Phi_{n+1}\rangle = b_{n+1}^2\langle\Phi_n|\Phi_n\rangle = \langle\Phi_n|H|\Phi_{n+1}\rangle. \quad (18)$$

The numerical implementation only requires keeping two N -dimensional vectors in memory. It is shown in table 1.

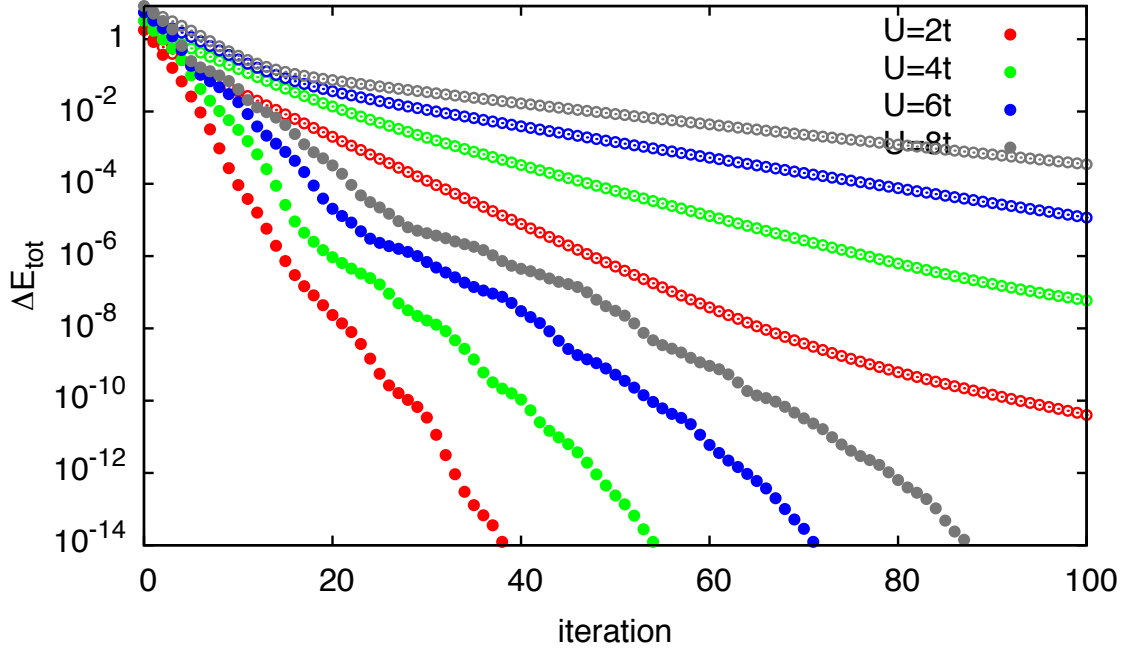


Fig. 2: *Covergence of the lowest eigenvalue for a Lanczos iteration (full circles) compared to steepest-descent minimization (open circles) of a 10-site Hubbard-chain at half-filling, starting from a random initial vector. Due to the additional variational degrees of freedom, Lanczos converges significantly faster. Overall, convergence for the half-filled system gets harder for larger U , as the distance to the lowest excited states is reduced ($\sim t^2/U$) and the spectrum widens ($\sim U$). In all cases, convergence is reached after less than $L \approx 100$ Lanczos iterations, to be compared to the dimension $N=63\,504$ of the Hilbert space.*

Diagonalizing (14), after a few tens to hundred iterations, the lowest eigenvalue of the tridiagonal representation of H on the Krylov space gives an excellent approximation to the ground-state energy of H in the full Hilbert space (Fig. 2). A formal estimate of the convergence was given by Kaniel and Paige [5]. For an $N+1$ -dimensional, symmetric matrix H with eigenvalues E_n , the lowest eigenvalue \check{E}_0 of the tridiagonal representation of H on the $(L+1)$ -dimensional Krylov space over $|v_0\rangle$ fulfills

$$\frac{\check{E}_0 - E_0}{E_N - E_0} \leq \left(\frac{\tan(\arccos(\langle \check{\Psi}_0 | \Psi_0 \rangle))}{T_L \left(1 + 2 \frac{E_1 - E_0}{E_N - E_1} \right)} \right)^2 \quad (19)$$

where $T_L(x)$ is the Chebyshev polynomial of order L and $\langle \check{\Psi}_0 | \Psi_0 \rangle$ the overlap of the Lanczos approximation to the ground-state $\check{\Psi}_0$ with the ground-state of H . Thus, if the initial state $|v_0\rangle$ is not orthogonal to the non-degenerate ground-state, convergence is exponential with a rate roughly increasing with the square root of the gap to the first excited measured in units of the width of the spectrum.

The approximate ground-state vector is given by the linear combination

$$|\check{\Psi}_0\rangle = \sum_{n=0}^L \check{\psi}_{0,n} |v_n\rangle, \quad (20)$$

where $\check{\psi}_0$ is the ground-state vector of the $L + 1$ -dimensional tridiagonal matrix (14). Instead of storing all $L + 1$ basis vectors $|v_n\rangle$, we can restart the Lanczos iteration from the same $|v_0\rangle$, accumulating the sum (20) iteration by iteration. This only requires keeping one additional N -dimensional vector in memory.

So far we have tacitly assumed that the Krylov vectors $H^n|v_0\rangle$ are linearly independent. If not, there will be a vector $H|\tilde{v}_m\rangle$ that vanishes when orthogonalized to the previous states, i.e., $b_m = 0$. This means that the Krylov space $\text{span}(|v_0\rangle, |v_1\rangle, \dots, |v_m\rangle)$ is invariant under H , i.e., we have found an exact eigenspace of H . For a large matrix H it is quite unlikely to be that lucky. Still, as the Lanczos iteration approaches the ground-state, we encounter a similar situation: Close to an eigenstate, the functional (1) becomes almost stationary, i.e., the coefficients b_n almost vanish. Normalization of the very short vector $|\tilde{v}_n\rangle$ then amplifies numerical noise in the small vector. This makes the numerical $|v_n\rangle$, which in theory should automatically be orthogonal to all $|v_m\rangle$ with $m < n - 2$, actually have finite overlaps with these vectors. This loss of orthogonality manifests itself in the appearance of multiple copies of eigenvectors (ghost states) which are unrelated to the actual multiplicities of the eigenvalues. This is the problem, which makes the Lanczos method unpractical for tridiagonalizing dense matrices. For the ground-state the variational principle prevents severe problem from the loss of orthogonality. An example of the appearance of ghost states is shown in figure 3.

If we want to reliably obtain excited states, we need to explicitly orthogonalize to the previous basis states. This leads to the Lanczos method with complete reorthogonalization [5]. A similar orthogonalization is performed in the Arnoldi method [2], which, however, is devised for unsymmetric matrices.

1.2 Spectral functions

Given the orthogonality problems of the Lanczos method, it appears hopeless to use it to obtain matrix elements of the resolvent, as they contain information about the full spectrum $H|\Psi_n\rangle = E_n|\Psi_n\rangle$. Still we are tempted to approximate the Lehmann representation

$$G_c(z) = \left\langle \Psi_c \left| \frac{1}{z - H} \right| \Psi_c \right\rangle = \sum_{n=0}^N \frac{\langle \Psi_c | \Psi_n \rangle \langle \Psi_n | \Psi_c \rangle}{z - E_n} \quad (21)$$

in terms of the eigenstates on the Krylov space $\mathcal{K}^L(|\Psi_c\rangle)$

$$\check{G}_c(z) = \left\langle \Psi_c \left| \frac{1}{z - \check{H}_c} \right| \Psi_c \right\rangle = \sum_{n=0}^L \frac{\langle \Psi_c | \check{\Psi}_n \rangle \langle \check{\Psi}_n | \Psi_c \rangle}{z - \check{E}_n}. \quad (22)$$

This is straightforward to calculate: We run L Lanczos iterations, starting from the (normalized) vector $|\Psi_c\rangle$, to create the tridiagonal \check{H}_c . The matrix element of the resolvent is the top left

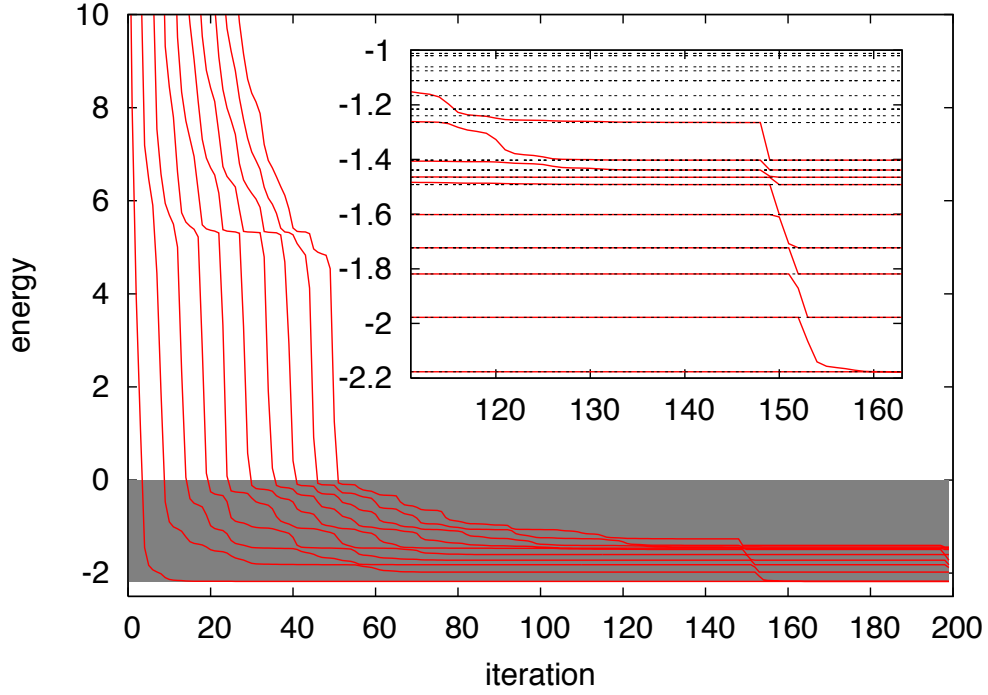


Fig. 3: Appearance of ghost states in an overconverged Lanczos iteration. The ground-state energy for a half-filled 8-site Hubbard-chain with $U = 10t$ is converged to numerical accuracy (10^{-16}) after about 85 iterations. Forcing the Lanczos iteration to continue, we see that at first also the higher excited states converge to the exact eigenvalues (dashed lines). But, as shown in the inset, they eventually start collapsing to the ground state. The appearance of these ghost states is due to the orthogonality problem introduced by small normalization parameters b_n , when the iteration is very close to a stationary point.

matrix element of the inverse of

$$z - \check{H}_c = \left(\begin{array}{c|ccccccc} z - a_0 & -b_1 & 0 & 0 & \cdots & 0 & 0 \\ \hline -b_1 & z - a_1 & -b_2 & 0 & \cdots & 0 & 0 \\ 0 & -b_2 & z - a_2 & -b_3 & \cdots & 0 & 0 \\ 0 & 0 & -b_3 & z - a_3 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & z - a_{L-1} & -b_L \\ 0 & 0 & 0 & 0 & \cdots & -b_L & z - a_L \end{array} \right). \quad (23)$$

This is easily determined, partitioning the matrix as indicated

$$z - \check{H}_c = \begin{pmatrix} z - a_0 & B^{(1)T} \\ B^{(1)} & z - \check{H}_c^{(1)} \end{pmatrix} \quad (24)$$

and inverting the block-matrix, giving

$$[(z - \check{H}_c)^{-1}]_{00} = \left(z - a_0 - B^{(1)T} (z - \check{H}_c^{(1)})^{-1} B^{(1)} \right)^{-1} = \left(z - a_0 - b_1^2 [(z - \check{H}_c^{(1)})^{-1}]_{00} \right)^{-1}.$$

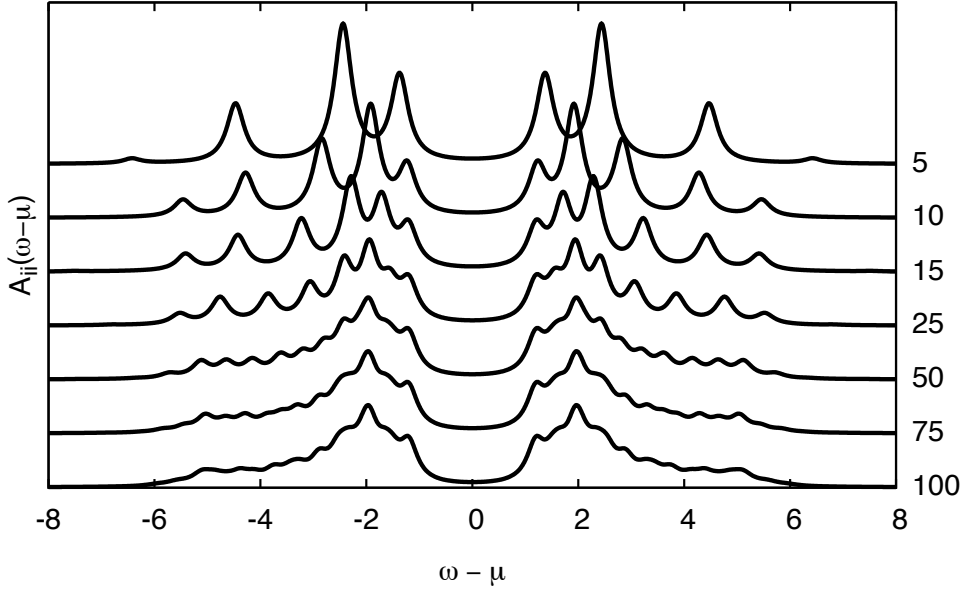


Fig. 4: Convergence of the spectral function with increasing number of Lanczos steps, $L=5, 10, 15, 25, 50, 75,$ and 100 , for a 14-site Hubbard chain with $U = 5t$ at half filling. With increasing L , more and more moments of the photoemission and inverse photoemission part of the spectrum are reproduced correctly.

Repeating inversion by partitioning for the submatrices $\check{H}^{(n)}$ we obtain the continued fraction

$$\check{G}_c(z) = [(z - \check{H}_c)^{-1}]_{00} = \frac{1}{z - a_0 - \frac{b_1^2}{z - a_1 - \frac{b_2^2}{z - a_2 - \dots}}}, \quad (25)$$

which terminates with $-b_L^2/(z - a_L)$. The spectral representation (22) is obtained by diagonalizing the Lanczos matrix \check{H}_c giving us the $L + 1$ eigenvalues \check{E}_n and eigenvectors $\check{\psi}_n$. Since

$$|\check{\Psi}_n\rangle = \sum_{l=0}^L \check{\psi}_{n,l} |v_l\rangle \quad (26)$$

the matrix elements are given by $\langle \check{\Psi}_n | \check{\Psi}_c \rangle = \check{\psi}_{n,0}$. Thus

$$\check{G}_c(z) = \sum_{n=0}^L \frac{|\check{\psi}_{n,0}|^2}{z - \check{E}_n} \quad (27)$$

The spectral function

$$\check{A}(\omega \pm i\eta) = \mp \frac{1}{\pi} \Im \check{G}(\omega \pm i\eta) \quad (28)$$

obtained this way converges very quickly. An example is shown in figure 4.

To understand how the $L + 1$ eigenstates of \check{H} can represent the full spectrum so well, we consider the moments of the spectral function

$$\int_{-\infty}^{\infty} d\omega \omega^m \check{A}(\omega) = \sum_{n=0}^L |\check{\psi}_{n,0}|^2 \check{E}_n^m = \sum_{n=0}^L \langle \check{\Psi}_c | \check{\Psi}_n \rangle \langle \check{\Psi}_n | \check{\Psi}_c \rangle \check{E}_n^m = \langle \check{\Psi}_c | \check{H}^m | \check{\Psi}_c \rangle \quad (29)$$

Since \check{H} is the projection of H onto the Krylov space $\mathcal{K}^L(|\check{\Psi}_c\rangle)$, we have $\check{H}^m |\check{\Psi}_c\rangle = H^m |\check{\Psi}_c\rangle$ for $m \leq L$. Thus the Lanczos representation $\check{A}(z)$ correctly reproduces the first $2L + 1$ moments of the spectral function $A(z)$. A further Lanczos step adds one new level to the continued fraction (25), leaving all previous terms unchanged. $b_m^2 = 0$ then implies that the continued fraction terminates, and all moments are given correctly. A near vanishing $b_m^2 \approx 0$, which gives rise to the loss of orthogonality of the Lanczos vectors, for the spectral function merely means that further terms in the continued fraction hardly contribute any more.

So far we have only considered diagonal elements of the resolvent. Off-diagonal matrix elements

$$G_{c_1, c_2}(z) = \left\langle \check{\Psi}_{c_2} \left| \frac{1}{z - H} \right| \check{\Psi}_{c_1} \right\rangle \quad (30)$$

are easily obtained by considering the diagonal elements for the linear combinations

$$\left\langle \check{\Psi}_{c_1} \pm \check{\Psi}_{c_2} \left| \frac{1}{z - H} \right| \check{\Psi}_{c_1} \pm \check{\Psi}_{c_2} \right\rangle = G_{c_1, c_1}(z) \pm G_{c_1, c_2}(z) \pm G_{c_2, c_1}(z) + G_{c_2, c_2}(z). \quad (31)$$

2 Application to the Hubbard model

The Hubbard model

$$H = -t \sum_{\langle i, j \rangle \sigma} c_{i\sigma}^\dagger c_{j\sigma} + U \sum n_{i\uparrow} n_{i\downarrow} \quad (32)$$

describes the fundamental dichotomy between itinerancy and locality for correlated electrons on a lattice: the hopping tends to delocalize electrons and is diagonal in k -space. This makes it possible to solve the band-structure problem for the infinite solid. In k -space the single electron Hamiltonian is block-diagonal. For the one-band Hubbard model each block is just the band energy ε_k . In general, each block defines the band structure problem for one k -point. Including electron-electron repulsion destroys this symmetry. The two-body Coulomb term is diagonal in real space, while in k -space it is dense

$$H = \sum_{k\sigma} \varepsilon_k c_{k\sigma}^\dagger c_{k\sigma} + \frac{U}{M} \sum_{k, k', q} c_{k\uparrow}^\dagger c_{k-q, \uparrow} c_{k'\downarrow}^\dagger c_{k'+q, \downarrow}. \quad (33)$$

This has two important consequences:

1. Since we know no general approach to transform the full Hamiltonian into finite-dimensional blocks, we have to restrict ourselves to finite-dimensional systems. For a cluster

M	N_\uparrow	N_\downarrow	dim of Hilbert space	memory
2	1	1	4	
4	2	2	36	
6	3	3	400	
8	4	4	4 900	
10	5	5	63 504	
12	6	6	853 776	6 MB
14	7	7	11 778 624	89 MB
16	8	8	165 636 900	1263 MB
18	9	9	2 363 904 400	17 GB
20	10	10	34 134 779 536	254 GB
22	11	11	497 634 306 624	3708 GB
24	12	12	7 312 459 672 336	53 TB
20	1	1	400	
20	2	2	36 100	
20	3	3	1 299 600	9 MB
20	4	4	23 474 025	179 MB
20	5	5	240 374 016	1833 MB
20	6	6	1 502 337 600	11 GB
20	7	7	6 009 350 400	44 GB
20	8	8	15 868 440 900	118 GB
20	9	9	28 210 561 600	210 GB
20	10	10	34 134 779 536	254 GB

Table 2: Dimension of Hilbert space $\dim(\mathcal{H})$ and computer memory required for storing a single many-body wave-function for Hubbard models with M orbitals and $N_\uparrow + N_\downarrow$ electrons. The first group of numbers gives the dimensions for half-filling, where the Hilbert space is largest. The second group shows how the dimension grows with the filling (dimensions are symmetric about half-filling). Note that the fourth column resembles a semi-logarithmic plot of $\dim(\mathcal{H})$ as a function of system size or filling.

of M sites with N_\uparrow electrons with spin up and N_\downarrow with spin down, the dimension of the Hilbert space is

$$\dim(\mathcal{H}) = \dim(\mathcal{H}_\uparrow) \times \dim(\mathcal{H}_\downarrow) = \binom{M}{N_\uparrow} \times \binom{M}{N_\downarrow}. \quad (34)$$

Examples for the single-band Hubbard model, illustrating the enormous growth of the Hilbert space are given in table 2. Actual calculations are therefore limited to quite small systems.

2. For a tight-binding system where hopping matrix-elements are restricted to close neighbors the many-body Hamiltonian is a sparse matrix when expressed in a real-space basis of localized orbitals (32). The basis states are then configurations $|\{n_{i\sigma}\}\rangle = \prod (c_{i\sigma}^\dagger)^{n_{i\sigma}} |0\rangle$, characterized by their occupation numbers $\{n_{i\sigma}\}$.

0	—	↑	↑↓	(0,0)
1	↑	—	↑↓	(0,1)
2	↑	↑	↓	(0,2)
3	—	↑↓	↑	(1,0)
4	↑	↓	↑	(1,1)
5	↑	↑↓	—	(1,2)
6	↓	↑	↑	(2,0)
7	↑↓	—	↑	(2,1)
8	↑↓	↑	—	(2,2)

Fig. 5: Basis configurations for a three site system with two up and one down spin electron. The left label denotes the index of the configuration. Equivalently, a state is also unambiguously pointed to by a tuple of up- and down-configuration index.

2.1 Representation of basis and Hamiltonian

Since the many-body basis states

$$|\{n_{i\sigma}\}\rangle = \prod_{i\sigma} \left(c_{i\sigma}^\dagger\right)^{n_{i\sigma}} |0\rangle \quad (35)$$

can be represented by Fermionic occupation numbers, it is natural to encode them in a string of bits. For a Hamiltonian like (32) that conserves spin, we can write

$$|\{n_{i\sigma}\}\rangle = \prod_{i=0}^{L-1} \left(c_{i\downarrow}^\dagger\right)^{n_{i\downarrow}} \left(c_{i\uparrow}^\dagger\right)^{n_{i\uparrow}} |0\rangle \quad (36)$$

with $\sum n_{i\sigma} = N_\sigma$, and encode each spin-component as the integer $m_\sigma = \sum n_{i\sigma} 2^i$. Enumerating all basis states with N_σ electrons on L sites is then as simple as looping over all integers from 0 to $2^L - 1$ and storing each integer m_σ with N_σ bits set. For $N_\uparrow = 2$ and $N_\downarrow = 1$ electrons on $L = 3$ sites we obtain

m_\uparrow	bits	state	i_\uparrow	m_\downarrow	bits	state	i_\downarrow
0	000			0	000		
1	001			1	001	$c_{0\downarrow}^\dagger 0\rangle$	0
2	010			2	010	$c_{1\downarrow}^\dagger 0\rangle$	1
3	011	$c_{0\uparrow}^\dagger c_{1\uparrow}^\dagger 0\rangle$	0	3	011		
4	100			4	100	$c_{2\downarrow}^\dagger 0\rangle$	2
5	101	$c_{0\uparrow}^\dagger c_{2\uparrow}^\dagger 0\rangle$	1	5	101		
6	110	$c_{1\uparrow}^\dagger c_{2\uparrow}^\dagger 0\rangle$	2	6	110		
7	111			7	111		

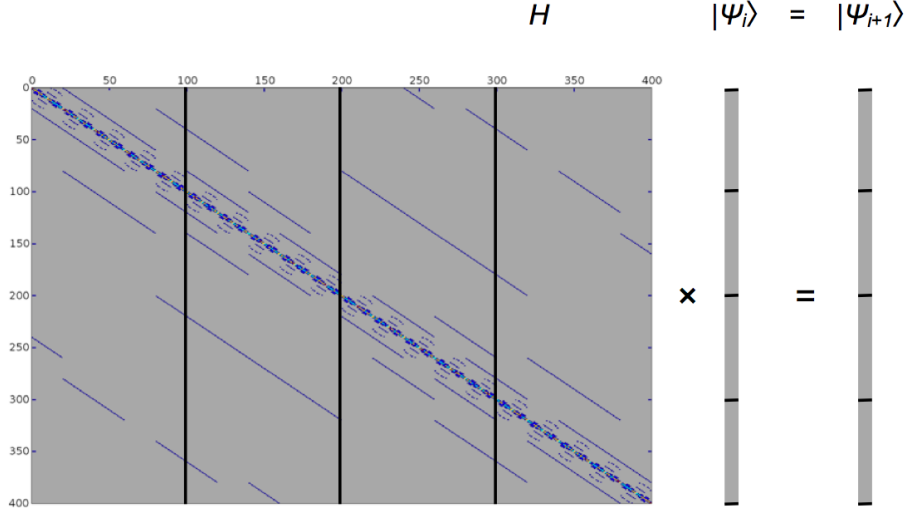


Fig. 6: Matrix-vector product for the Lanczos iteration showing (in blue) the non-zero elements of the Hamiltonian matrix for a six-site Hubbard model at half-filling. Vector Ψ_i is only read, Ψ_{i+1} can be written sequentially: $\Psi_{i+1,n} = \sum_m H_{n,m} \Psi_{i,m}$. Access to elements of Ψ_i is highly non-local.

We number the basis states in the order they are found $i_\sigma = 0 \dots \dim(\mathcal{H}_\uparrow)$. A full basis state (36) is then indexed by the $i = i_\downarrow + \dim(\mathcal{H}_\downarrow) \cdot i_\uparrow$. This corresponds to writing the basis as the tensor product of the up- and down-states. Alternatively we could use $i = i_\uparrow + \dim(\mathcal{H}_\uparrow) \cdot i_\downarrow$. The corresponding configuration is given by the integers $m_\uparrow[i_\uparrow]$ and $m_\downarrow[i_\downarrow]$

For calculating the matrix elements it is convenient to store the two lookup tables for converting between the integer encoding the basis state m_σ and its index i_σ in the basis. Since $\dim(\mathcal{H}_\sigma)$ is normally (i.e. close to half-filling) much smaller than $\dim(\mathcal{H})$, this does not use much memory.

The hopping term connects basis state that differ only in two occupation numbers of the same spin, e.g., $n_{i\sigma}$ and $n_{j\sigma}$. The matrix element is $\mp t_{ij}$, where the sign depends on the number $N_{i,j}$ of electrons of spin σ between site i and site j : $-(-1)^{N_{i,j}} t_{i,j}$. As an example we give the matrix of the hopping between the basis states for $N_\uparrow = 2$ electrons on a linear cluster with $L = 3$ sites, nearest neighbor hopping t and with periodic boundary conditions:

$$T_\uparrow = \begin{pmatrix} 0 & -t & +t \\ -t & 0 & -t \\ +t & -t & 0 \end{pmatrix}. \quad (37)$$

The full hopping matrix is then given by the tensor product of T_\uparrow and T_\downarrow . The matrix above looks fairly dense, for larger system the T quickly becomes very sparse, as shown in figure 6.

2.2 Green functions

In a basis of spin-orbitals α and β , the elements of the Green matrix are given by

$$\begin{aligned}
 G_{\alpha\beta}(\omega) &= \left\langle \Psi_0 \left| c_\alpha^\dagger \frac{1}{\omega + (H - E_0 - i\eta)} c_\beta \right| \Psi_0 \right\rangle + \left\langle \Psi_0 \left| c_\alpha \frac{1}{\omega - (H - E_0 - i\eta)} c_\beta^\dagger \right| \Psi_0 \right\rangle \\
 &= \sum_n \frac{\langle \Psi_0 | c_\alpha^\dagger | \Psi_n^{(N-1)} \rangle \langle \Psi_n^{(N-1)} | c_\beta | \Psi_0 \rangle}{\omega + (E_n^{(N-1)} - E_0^{(N)}) - i\eta} + \sum_n \frac{\langle \Psi_0 | c_\alpha | \Psi_n^{(N+1)} \rangle \langle \Psi_n^{(N+1)} | c_\beta^\dagger | \Psi_0 \rangle}{\omega - (E_n^{(N+1)} - E_0^{(N)}) + i\eta},
 \end{aligned} \tag{38}$$

where the sums are over the eigenstates of the Hilbert space with one electron less (first term) and one additional electron (second term). Diagonal elements are calculated in Lanczos as described in section 1.2: To find $G_{\alpha\alpha}(\omega)$, we need the ground state vector $|\Psi_0\rangle$ and two additional Lanczos runs, giving the two terms in (38). For the first term, we start the Lanczos iteration from the normalized vector $|\Psi_c^<\rangle = c_\alpha |\Psi_0\rangle / \sqrt{n_\alpha}$, where $n_\alpha = \langle \Psi_0 | c_\alpha^\dagger c_\alpha | \Psi_0 \rangle$. Likewise, for the second term, we start from $|\Psi_c^>\rangle = c_\alpha^\dagger |\Psi_0\rangle / \sqrt{1 - n_\alpha}$. The Green function is then given, in terms of the Lanczos coefficients, by

$$\check{G}_{\alpha\alpha}(\omega) = \frac{n_\alpha}{\omega - E_0 - i\eta + a_0^< - \frac{b_1^{<2}}{\omega - E_0 - i\eta + a_1^< - \dots}} + \frac{1 - n_\alpha}{\omega + E_0 + i\eta - a_0^> - \frac{b_1^{>2}}{\omega + E_0 + i\eta - a_1^> - \dots}}. \tag{39}$$

If the ground state is degenerate, e.g., for $N_\uparrow \neq N_\downarrow$, where $E_0(N_\uparrow, N_\downarrow) = E_0(N_\downarrow, N_\uparrow)$, we average the Green functions calculated from the different ground-states. This is the $T \rightarrow 0$ limit of the finite-temperature Green function

$$G_{\alpha\alpha}(\omega) = \frac{1}{Z} \sum_m e^{-\beta E_m^{(N)}} G_{\alpha\alpha}^{(m)}(\omega), \tag{40}$$

where $Z = \sum_n e^{-\beta E_n^{(N)}}$ is the partition function and $G_{\alpha\alpha}^{(m)}(\omega)$ has the same form as (38), only with Ψ_0 replaced by Ψ_m . For finite, but sufficiently low temperatures, the Boltzmann factor is negligibly small, except for the lowest few states. If we calculate those, taking care of the orthogonality problem (ghost states), we can easily obtain the finite-temperature Green function. A more elaborate method is given in [12].

Off-diagonal elements of the Green matrix are calculated from diagonal elements of linear combinations of spin-orbitals, e.g., $(c_\alpha^\dagger \pm c_\beta^\dagger) |\Psi_0\rangle$, as described in section 1.2.

2.3 Parallelization strategies

Because of the enormous size of the many-body Hilbert space, see table 2, Lanczos calculations are limited by the available memory. On shared-memory systems the most time consuming operation of the Lanczos iteration, the multiplication of the Hamiltonian-matrix with a many-body vector, can be parallelized very easily, when it is written such that the elements of the resulting vector are calculated independently: As illustrated in figure 6, different threads can

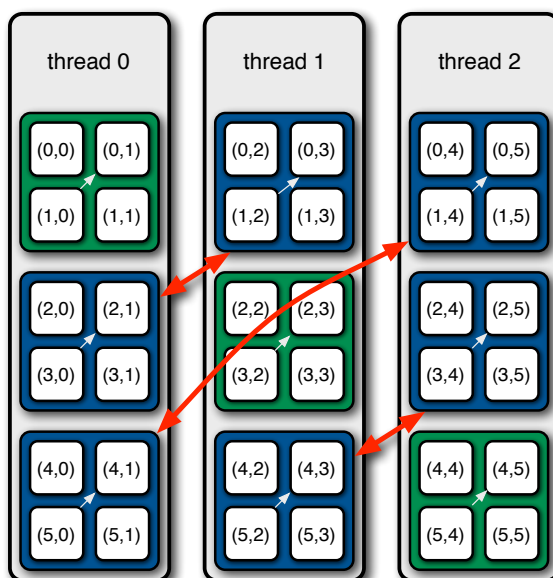


Fig. 7: Transpose operation that makes memory access thread-local when calculating the operation of the Hamiltonian on the state-vector. The communication (red arrows) is realized by a call to `MPI_Alltoall`, which is very efficiently implemented on Blue Gene/L. The small grey arrows indicate the local operations needed to complete the matrix-transpose.

work on different chunks of $|\Psi_{i+1}\rangle$. The off-diagonal elements of the kinetic energy part of (32) lead to non-local memory access, but the elements of $|\Psi_i\rangle$ as well as the matrix elements are only read, so that there is no need for locking. An OpenMP parallelization thus needs only a single pragma. Parallelizing also the scalar products in a similar way, we obtain almost ideal speedup on an IBM p690 frame of JUMP in Jülich. The implementation is however limited to a single node, i.e., about 120 GBytes. To use significantly more memory we need to find an efficient distributed-memory implementation.

A naive approach on distributed memory systems uses MPI2 one-sided communication to emulate the shared-memory approach by direct remote memory access. This leads, however, to a severe speed-down, i.e., the more processors we use, the longer we have to wait for the result.

An efficient distributed-memory implementation [14] is instead based on the fact that hopping does not change spin. Hopping of the up-electron mixes only different up-hopping configurations, while the down-electron configuration remains unchanged. If we group all up configurations for a fixed down configuration together in a single thread, this hopping can be carried out locally. Figure 5 illustrates this: for a fixed index i_\downarrow , all i_\uparrow configurations are stored in adjacent memory locations and can be stored in a thread. We see, that this basis can be naturally indexed by a tuple $(i_\downarrow, i_\uparrow)$ (right labels in figure 5) instead of the global index (left labels). We can therefore equivalently regard the vectors as matrices $v(i_\downarrow, i_\uparrow)$ with indices i_\downarrow and i_\uparrow . Now it is easy to see that a matrix transpose reshuffles the data elements such that the down configurations are sequentially in memory and local to the thread.

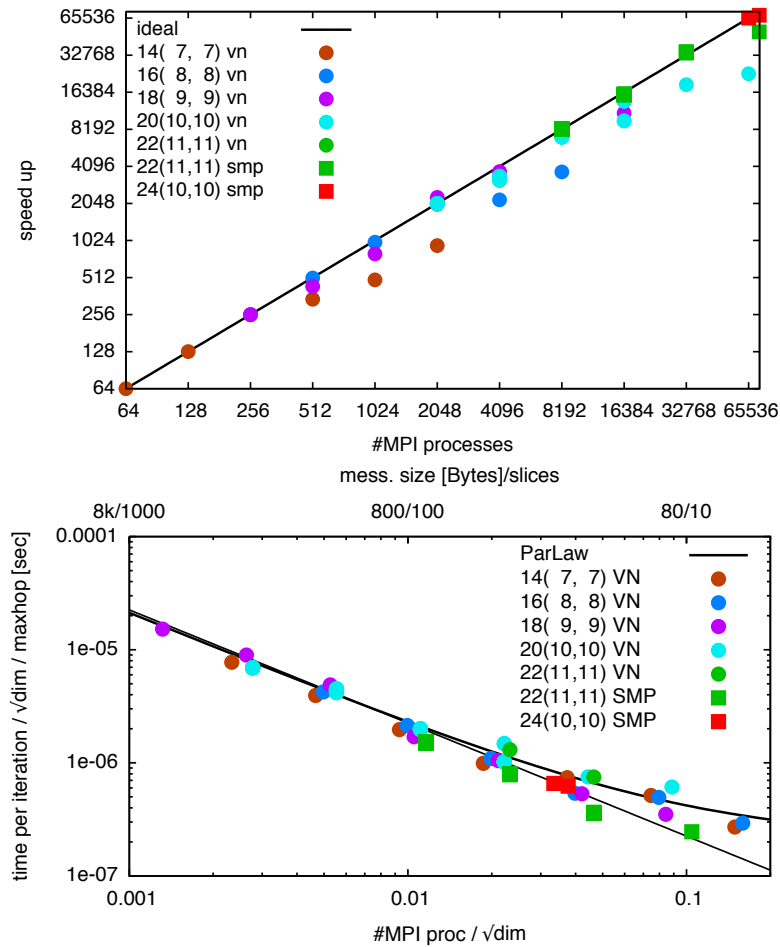


Fig. 8: Timings of the parallel implementation of the Lanczos algorithm for the Hubbard model on the Jülich IBM BlueGene/P. Sizes of the state vectors of the half-filled systems are given in table 2. For the 24 site system with 10+10 electrons, $\dim(\mathcal{H}) = 3\,846\,525\,097\,536$, the state vector takes about 28 TBytes. The simulation of such a system requires the entire machine, using only one processor per node (SMP mode), to make most efficient use of the available memory. For smaller systems we can use all four processors per node (VN mode). Despite massive communication in each iteration, the code shows excellent speed up. Only when the message size per process become too small, performance degrades because of network latency. This is shown in the lower plot. Properly scaling the execution times we obtain a universal scaling (ParLaw) for system sizes ranging over more than five and process counts ranging over three orders of magnitude.

We implement an efficient matrix transpose using `MPI_Alltoall`. This routine expects that the data packages which will be sent to a given process to be stored contiguously in memory. This does not apply to our case, since we would like to store the spin-down electron configurations sequentially in memory. Thus, the matrix is stored column wise. For `MPI_Alltoall` to work properly, we would have to bring the data elements into row-major order. This could be done by performing a local matrix transpose. The involved matrices are, however, in general rectangular, leading to expensive local-copy and reordering operations. We can avoid this by calling `MPI_Alltoall` for each column separately (red arrows in figure 7). After this, only

a local strided transposition has to be performed (small white arrows) to obtain the fully transposed matrix or Lanczos vector. The implementation described so far uses `MPI_Alltoall` which assumes that the matrix to be transposed is a square matrix and that the dimension $dim_{\uparrow} = dim_{\downarrow}$ is divisible by the number of MPI processes. To overcome these restrictions we have generalized the algorithm to `MPI_Alltoallv`. This is the implementation that is used in practice. The speed-up shown at the top of figure 8 shows that our parallelization based on collective communication is indeed very efficient. Even for a system of 24 sites with 10 electrons of either spin, where a single many-body vector takes about 28 TB of memory, our implementation works very efficiently despite the fact that in each Lanczos iteration 28 TB of data have to be moved across the entire machine twice.

The lower plot in figure 8 shows that the execution times for runs of systems varying by more than five orders of magnitude in size (of the Hilbert space) for processor counts varying over three orders of magnitude fall on a universal curve, which is only determined by the bandwidth and the latency of the network. This suggests that the implementation should scale to even larger systems than the present Jülich BlueGene/P with almost 300 000 CPUs and an aggregate memory of 144 terabytes.

3 Application to DMFT

Using the Lanczos method as a solver for DMFT gives results at zero temperature and directly on the real axis. An important limitation is, however, the need to approximate the bath Green's function

$$\mathcal{G}^{-1}(\omega) = \omega + \mu - \int_{-\infty}^{\infty} d\omega' \frac{\Delta(\omega')}{\omega - \omega'} \quad (41)$$

by a discretized version, e.g., of the form

$$\mathcal{G}_{\text{And}}^{-1}(\omega) = \omega + \mu - \sum_{l=1}^{N_b} \frac{V_l^2}{\omega - \varepsilon_l}, \quad (42)$$

corresponding to an Anderson impurity model with a finite number of sites

$$H_{\text{And}} = \varepsilon_0 \sum_{\sigma} n_{\sigma} + U n_{\uparrow} n_{\downarrow} + \sum_{\sigma} \sum_{l=1}^{N_b} \left(\varepsilon_l n_{l\sigma} + V_l \left(a_{l\sigma}^{\dagger} c_{\sigma} + c_{\sigma}^{\dagger} a_{l\sigma} \right) \right), \quad (43)$$

where c_{σ}^{\dagger} and $a_{l\sigma}^{\dagger}$ create an electron of spin σ on the impurity or bath-site l , respectively, $n_{\sigma} = c_{\sigma}^{\dagger} c_{\sigma}$ and $n_{l\sigma} = a_{l\sigma}^{\dagger} a_{l\sigma}$. Writing the non-interacting part of H_{And} as a matrix

$$H_{\text{And}}^0 = \begin{pmatrix} 0 & V_1 & V_2 & V_3 & \cdots \\ V_1 & \varepsilon_1 & 0 & 0 & \\ V_2 & 0 & \varepsilon_2 & 0 & \\ V_3 & 0 & 0 & \varepsilon_3 & \\ \vdots & & & & \ddots \end{pmatrix} \quad (44)$$

we see that (42) is easily recovered from inversion by partitioning.

Since practical calculations are limited by the rapidly increasing size of the Hilbert space to small numbers of bath-sites N_b , it is crucial for the reliability of the calculations to find a good representation $\mathcal{G}_{\text{And}}^{-1}$ for the bath Green's function. The most common approach is to use a least squares fit [15]: Because of the spectral poles on the real axis, such a fit in practice is done on the imaginary axis, where the Green's functions are smooth and the optimization of the distance function is not easily trapped in local minima. One then minimizes a function of the form

$$\chi^2(\{V_l, \varepsilon_l\}) = \sum_{n=0}^{n_{\max}} |\mathcal{G}^{-1}(i\omega_n) - \mathcal{G}_{\text{And}}^{-1}(i\omega_n)|^2 \quad (45)$$

on a set of Matsubara frequencies, corresponding to some fictitious temperature. The choice of this temperature and of n_{\max} essentially determines the relative weighing of high versus low frequency features in the fit. If low $i\omega_n$ are weighted too little, the fit easily becomes underdetermined, since for large imaginary frequencies the hybridization function contains only little information about the system (which is the reason why the analytic continuation back to the real axis is so difficult). To emphasize different frequency ranges, it is possible to introduce frequency-dependent weight functions in (45).

Instead of fitting, we could use a moment expansion of the Weiss function $W(\omega) = \int d\omega' \frac{\Delta(\omega')}{\omega - \omega'}$ similar to that discussed in section 1.2. Such an approach [16] has been used for the Bethe lattice with infinite coordination, where the self-consistency condition simplifies to $W(\omega) = t^2 G_{\text{imp}}(\omega)$: As Lanczos gives a continued-fraction representation for the photoemission and inverse-photoemission part separately, the hybridization function is written as

$$W^<(\omega) + W^>(\omega) = t^2 G^<(\omega) + t^2 G^>(\omega) = \frac{t^2 b_0^{<}{}^2}{\omega + a_0^<} - \frac{b_1^{<}{}^2}{\omega + a_1^<} - \dots + \frac{t^2 b_0^{>}{}^2}{\omega - a_0^>} - \frac{b_1^{>}{}^2}{\omega - a_1^>} - \dots \quad (46)$$

Truncating the continued fractions at $N_b^<$ and $N_b^>$, this corresponds to the impurity model with

$$H_{\text{And}}^0 = \begin{pmatrix} 0 & t^2 b_0^< & & \dots & t^2 b_0^> \\ t^2 b_0^< & -a_0^< & b_1^< & & \\ & b_1^< & -a_1^< & b_2^< & \\ & & b_2^< & -a_2^< & \ddots \\ \vdots & & & \ddots & \ddots \\ t^2 b_0^> & & & & a_0^> & b_1^> \\ & & & & b_1^> & a_1^> & b_2^> \\ & & & & & b_2^> & a_2^> & \ddots \\ & & & & & & & \ddots \end{pmatrix}, \quad (47)$$

where the bath forms two chains, coupled to the impurity. Diagonalizing the bath, it is easily brought to the form (43). This bath-parametrization works very well for systems with large gap. An example is shown in figure 9. Since the approach uses moment expansions for the two parts

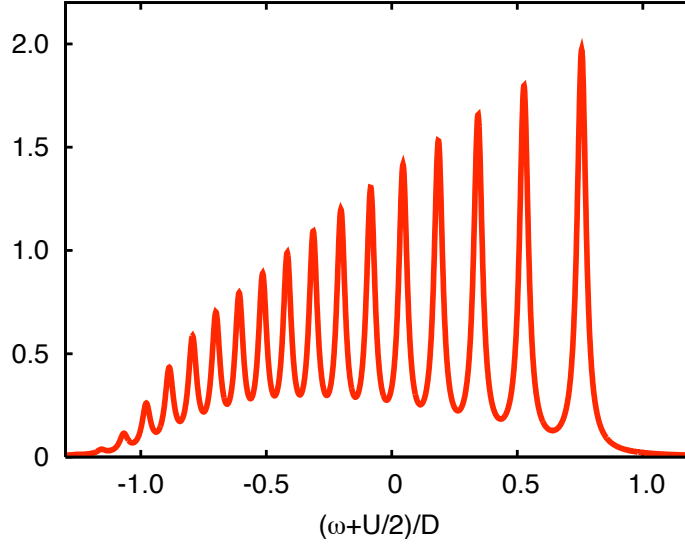


Fig. 9: Spectral function of the lower Hubbard band of a Hubbard model on the infinite Bethe lattice with half-bandwidth D and $U = 8D$ in antiferromagnetic DMFT. The bath was obtained from the continued-fraction expansion of the impurity Green's functions, $N_b = 24$ [18].

$W^<$ and $W^>$ of the Weiss function separately, it does not converge as quickly as a moment expansion for the full hybridization function would. This makes itself particularly felt, when the gap is small or the system is even metallic. To improve the description of the hybridization function close to the Fermi level, it has been proposed to introduce an additional bath site with fixed energy close to μ [17].

3.1 Cluster methods

For cluster versions of DMFT we can use exact sum-rules and symmetries to find the structure of the bath. Our discussion will closely follow [19]. To fix the notation we briefly sketch the self-consistency loop for cellular DMFT (CDMFT) and the dynamical cluster approximation (DCA) using Lanczos as impurity solver. Let N_c be the number of cluster-sites, N_b the number of bath-sites. For simplicity we suppress spin-indices.

Given an $N_c \times N_c$ bath Green matrix \mathcal{G}^{-1} ,

1. fit parameters of an Anderson model with N_b bath-sites

$$\mathcal{G}_{\text{And}}^{-1}(\omega) = \omega + \mu - \mathbf{H}_c - \mathbf{\Gamma} [\omega - \mathbf{E}]^{-1} \mathbf{\Gamma}^\dagger \quad (48)$$

to \mathcal{G}^{-1} , where $\mathbf{\Gamma}$ is the $N_c \times N_b$ -dimensional hybridization matrix, and \mathbf{E} the $N_b \times N_b$ -dimensional bath-matrix. \mathbf{H}_c is specified below,

2. solve the $(N_c + N_b)$ -site Anderson model H_{And} (specified below) to obtain the $N_c \times N_c$ cluster Green matrix \mathbf{G}_c ,

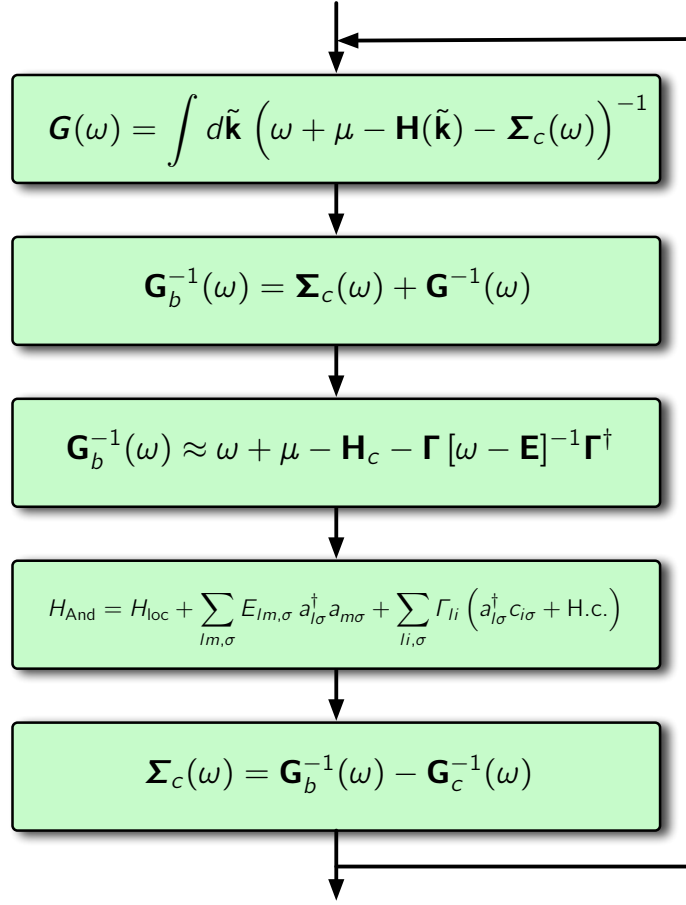


Fig. 10: Self-consistency loop for (cluster) DMFT.

3. get the cluster self-energy matrix

$$\boldsymbol{\Sigma}_c(\omega) = \mathcal{G}^{-1}(\omega) - \mathbf{G}_c^{-1}(\omega), \quad (49)$$

4. calculate the local Green matrix for the cluster by integrating over the reduced Brillouin-zone of the cluster

$$\mathbf{G}(\omega) = \int d\tilde{\mathbf{k}} \left(\omega + \mu - \mathbf{H}(\tilde{\mathbf{k}}) - \boldsymbol{\Sigma}_c(\omega) \right)^{-1}, \quad (50)$$

where $\mathbf{H}(\tilde{\mathbf{k}})$ is the single-electron part of the of the Hubbard Hamiltonian (32) in the reduced Brillouin-zone of the cluster,

5. determine the new bath Green matrix (self-consistency condition)

$$\mathcal{G}^{-1}(\omega) = \boldsymbol{\Sigma}_c(\omega) + \mathbf{G}^{-1}(\omega). \quad (51)$$

These steps are iterated to self-consistency.

3.2 Anderson impurity model

The Anderson model to be solved in step 2 is given by

$$H_{\text{And}} = H_{\text{clu}} + \sum_{lm,\sigma} E_{lm,\sigma} a_{l\sigma}^\dagger a_{m\sigma} + \sum_{li,\sigma} \Gamma_{il} \left(a_{l\sigma}^\dagger c_{i\sigma} + \text{H.c.} \right) \quad (52)$$

where the operator $a_{l\sigma}^\dagger$ creates an electron of spin σ on bath-site l . The cluster Hamiltonian H_{clu} is obtained from the lattice Hamiltonian by transforming to the reciprocal space of the super-lattice of the clusters, and projecting to the cluster. Writing the single-electron part of $H(\tilde{\mathbf{k}})$ as the matrix $\mathbf{H}(\tilde{\mathbf{k}})$, the single-electron part of H_{clu} is given by

$$\mathbf{H}_c = \int d\tilde{\mathbf{k}} \mathbf{H}(\tilde{\mathbf{k}}). \quad (53)$$

The (local) interaction terms are simply those of the lattice model, restricted to the cluster. The Hamiltonian $H(\tilde{\mathbf{k}})$ in the reciprocal space of the super-lattice $\{\tilde{\mathbf{r}}\}$ of clusters can be obtained by changing to the basis of operators

$$\tilde{c}_{\mathbf{R}_i\sigma}^{\text{CDMFT}}(\tilde{\mathbf{k}}) = \sum_{\tilde{\mathbf{r}}} e^{-i\tilde{\mathbf{k}}\tilde{\mathbf{r}}} c_{\tilde{\mathbf{r}}+\mathbf{R}_i,\sigma}. \quad (54)$$

The resulting quantum cluster approximation is CDMFT. Alternatively, we can start from the operators in the reciprocal space of the *lattice* to obtain

$$\tilde{c}_{\mathbf{R}_i\sigma}^{\text{DCA}}(\tilde{\mathbf{k}}) = \sum_{\tilde{\mathbf{r}}} e^{-i\tilde{\mathbf{k}}(\tilde{\mathbf{r}}+\mathbf{R}_i)} c_{\tilde{\mathbf{r}}+\mathbf{R}_i,\sigma}. \quad (55)$$

Now we obtain the DCA. The choice of the operators in the two approaches differs just by local phase factors. In CDMFT this Kohn-gauge [20] is chosen such that phases appear only in matrix elements involving different clusters. Thus all matrix elements on the cluster are the same as in the original Hamiltonian. The price for retaining the original matrix elements on the cluster is a breaking of the translation-symmetry of the original lattice. DCA opts instead to retain this symmetry by distributing the phase-change uniformly over the cluster-sites. The price for retaining translation-invariance is that the matrix elements in the cluster Hamiltonian differ from those in the original Hamiltonian (coarse graining). In both cases, CDMFT and DCA, the eigenvalues of $\mathbf{H}(\tilde{\mathbf{k}})$ are identical to the eigenvalues of the non-interacting part of H . Obviously, we can construct other cluster extensions to DMFT by different choices of the Kohn-gauge $\varphi(\tilde{\mathbf{k}}; \mathbf{R}_i)$ on the cluster

$$\tilde{c}_{\mathbf{R}_i\sigma}^\varphi(\tilde{\mathbf{k}}) = \sum_{\tilde{\mathbf{r}}} e^{-i(\tilde{\mathbf{k}}\tilde{\mathbf{r}}+\varphi(\tilde{\mathbf{k}};\mathbf{R}_i))} c_{\tilde{\mathbf{r}}+\mathbf{R}_i,\sigma}. \quad (56)$$

3.3 Hybridization sum-rules

While the most general parametrization for the bath is given by expression (48), we can always diagonalize the hopping matrix \mathbf{E} among the bath-sites to obtain

$$\mathcal{G}_{\text{And}}^{-1}(\{\varepsilon_l, \mathbf{V}_l\}; \omega) = \omega + \mu - \mathbf{H}_c - \sum_l \frac{\mathbf{V}_l \mathbf{V}_l^\dagger}{\omega - \varepsilon_l}. \quad (57)$$

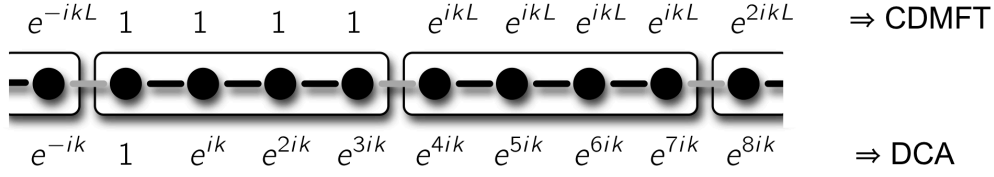


Fig. 11: Phase choice on the cluster that leads to CDMFT or DCA.

The hybridization matrix is then given by the tensor product of the vectors \mathbf{V}_l , where

$$V_{l,i} = \sum_m \Gamma_{i,m} \phi_{l,m} \quad (58)$$

and ϕ_l are the eigenvectors of \mathbf{E} with eigenvalues ε_l .

To obtain sum-rules for the hybridizations, we write the inverse of the bath Green matrix as

$$\mathcal{G}^{-1}(\omega) = \Sigma_c(\omega) + \left(\int d\tilde{\mathbf{k}} \left(\omega + \mu - \mathbf{H}(\tilde{\mathbf{k}}) - \Sigma_c(\omega) \right)^{-1} \right)^{-1}.$$

Considering the limit $\omega \rightarrow \infty$, expanding to order $1/\omega^2$, using (53), and comparing to (57) we find

$$\sum_l \mathbf{V}_l \mathbf{V}_l^\dagger = \int d\tilde{\mathbf{k}} \mathbf{H}^2(\tilde{\mathbf{k}}) - \left(\int d\tilde{\mathbf{k}} \mathbf{H}(\tilde{\mathbf{k}}) \right)^2. \quad (59)$$

To illustrate this hybridization sum-rule we consider a representative set of examples.

Single site

We consider a d -dimensional lattice with hoppings t_n to the z_n n^{th} -nearest neighbors. For $N_c = 1$ we have $\mathbf{H}(\mathbf{k}) = \varepsilon_{\mathbf{k}}$. Thus we find for the hybridizations

$$\sum_l V_l^2 = \frac{1}{(2\pi)^d} \int_{-\pi}^{\pi} d^d \mathbf{k} \varepsilon_{\mathbf{k}}^2 = \sum_n z_n t_n^2, \quad (60)$$

where the integral is just the second moment of the density of states, so that the last equation follows as in the recursion method [7]. For a Bethe lattice of connectivity z with hopping matrix element t/\sqrt{z} the sum-rule reduces to $\sum_l V_l^2 = t^2$.

CDMFT

We start by considering a linear chain with nearest neighbor hopping t and a three-site cluster $N_c = 3$. In the CDMFT gauge we have

$$\mathbf{H}(\tilde{k}) = -t \begin{pmatrix} 0 & 1 & e^{-3i\tilde{k}} \\ 1 & 0 & 1 \\ e^{3i\tilde{k}} & 1 & 0 \end{pmatrix} \quad (61)$$

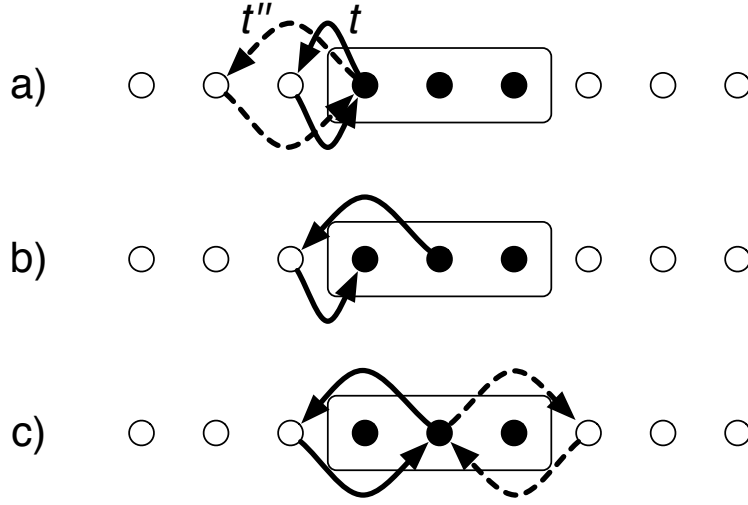


Fig. 12: CDMFT sum-rules for a one-dimensional 3-site cluster with nearest and next-nearest neighbor hoppings t and t'' , respectively: a) $\sum_l |V_{l,1}|^2 = t^2 + t''^2$, b) $\sum_l \bar{V}_{l,1} V_{l,2} = t t''$, and c) $\sum_l |V_{l,2}|^2 = 2t''^2$. The hybridizations are given by the two-step hopping processes that are lost when the cluster is cut out of the original lattice.

so that \mathbf{H}_c is the original single-electron Hamiltonian restricted to the cluster:

$$\mathbf{H}_c = \frac{3}{2\pi} \int_{-\pi/3}^{\pi/3} d\tilde{k} \mathbf{H}(\tilde{k}) = -t \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}. \quad (62)$$

The sum-rule (59) then is

$$\left(\sum_l V_{l,i} \bar{V}_{l,j} \right) = \begin{pmatrix} t^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & t^2 \end{pmatrix}, \quad (63)$$

i.e., only the sites on the surface of the cluster couple to the bath. If we allow also second nearest neighbor hopping with matrix element t'' , we find

$$\left(\sum_l \bar{V}_{l,\mu} V_{l,\nu} \right) = \begin{pmatrix} t^2 + t''^2 & t t'' & 0 \\ t t'' & 2t''^2 & t t'' \\ 0 & t t'' & t^2 + t''^2 \end{pmatrix}. \quad (64)$$

The general CDMFT hybridization sum-rule (59) can be easily visualized: The integral over the Brillouin-zone of the cluster projects the single-electron part of the full Hamiltonian to the cluster (see eqn. (53)). The matrix elements of \mathbf{H}_c^2 are thus the two-step hoppings that are possible on the cluster. Likewise the integral over the Hamiltonian squared gives the second moments, only that here the intermediate site is not restricted to the cluster. Thus the sum-rule matrix is given by the second-order paths between cluster-sites that proceed via a site outside the cluster. This is illustrated in figure 12. As a special case, for a single site we recover the second equality in (60).

The vanishing of a matrix element in the sum-rule merely implies that the corresponding matrix element of the bath Green matrix decays faster than $1/\omega$ for large ω . For a diagonal element, however, all terms in $\sum_l V_{l,i} \bar{V}_{l,i}$ are positive. Thus a vanishing sum means that all terms must be zero. Hence the sum-rule implies that cluster-sites that are so far in the interior that they cannot be reached by hopping from outside the cluster do not couple to bath and that all matrix elements of the bath Green function involving such a site i are given by $\mathcal{G}_{ij}^{-1}(\omega) = \omega + \mu - (\mathbf{H}_c)_{ij}$ for all ω . In that sense the bath hybridizes only to the surface of the cluster and we see that the hybridization-strength to these sites does not decrease for increasing cluster size N_c .

DCA

We start again by considering the 3-site cluster. In the DCA gauge we write

$$\mathbf{H}(\tilde{k}) = -t \begin{pmatrix} 0 & e^{i\tilde{k}} & e^{-i\tilde{k}} \\ e^{-i\tilde{k}} & 0 & e^{i\tilde{k}} \\ e^{i\tilde{k}} & e^{-i\tilde{k}} & 0 \end{pmatrix}. \quad (65)$$

Now \mathbf{H}_c has translation symmetry, but the hopping is rescaled by $\sin(\pi/N_c)/(\pi/N_c)$

$$\mathbf{H}_c = \frac{3}{2\pi} \int_{-\pi/3}^{\pi/3} d\tilde{k} \mathbf{H}(\tilde{k}) = -\frac{3\sqrt{3}}{2\pi} t \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}. \quad (66)$$

Since all matrices in (59) are periodic, it is convenient to transform to k -space. With $V_{l,K} = \sum_i V_{l,i} e^{iKr_i} / \sqrt{N_c}$ and the coarse-graining factor $\tau = 3\sqrt{3}/2\pi$ we find

$$\begin{aligned} \sum_l |V_{l,K=0}|^2 &= (2 + \tau - 4\tau^2) t^2 \\ \sum_l |V_{l,K=\pm 2\pi/3}|^2 &= (2 - \tau/2 - \tau^2) t^2. \end{aligned}$$

The hybridization sum-rule (59) is then, likewise, diagonal in the cluster-momenta \mathbf{K}

$$\sum_l |V_{l,\mathbf{K}}|^2 = \int d\tilde{\mathbf{k}} \varepsilon_{\mathbf{K}+\tilde{\mathbf{k}}}^2 - \left(\int d\tilde{\mathbf{k}} \varepsilon_{\mathbf{K}+\tilde{\mathbf{k}}} \right)^2, \quad (67)$$

while all terms $V_{l,\mathbf{K}} \bar{V}_{l,\mathbf{K}'}$ mixing different cluster momenta vanish. As a special case, for a single site the above sum-rule is just the first equality in (60). Expanding $\varepsilon_{\mathbf{K}+\tilde{\mathbf{k}}}$ around \mathbf{K} , we find that for a d -dimensional system $\sum_l |V_{l,\mathbf{K}}|^2$ decreases with cluster size as $1/N_c^{2/d}$, while all cluster-sites couple with the same strength to the bath.

Discussion

Besides providing exact relations for the bath-parametrization, in particular which sites need not be coupled to a bath, the sum-rules contain important information about the scaling of cluster methods with cluster size: In CDMFT individual hybridizations are independent of cluster size,

while for DCA they decrease with cluster size as $N_c^{-2/d}$. Interestingly this means that for a d -dimensional system in CDMFT the overall coupling to the bath scales as $N_c^{(d-1)/d}$, while in DCA it scales as $N_c^{(d-2)/d}$. For non-local properties a DCA calculation is therefore expected to converge faster with cluster size. For a calculation where we represent the bath by discrete degrees of freedom this decrease in hybridization strength does, however, not help very much as we still need bath-sites to fit the hybridizations, even if they are small. With increasing DCA cluster size we thus have to parametrize N_c baths, one for each \mathbf{K} . In CDMFT the situation is more fortunate, as the sum-rules imply that many hybridizations vanish and we only need to parametrize the coupling of surface-sites to the bath.

The lack of translational invariance in CDMFT has, however, two important practical implications. First, the full Green matrix has to be calculated, instead of just its diagonal. Second, when calculating local quantities, like the density per site, in CDMFT we have a choice of inequivalent sites, or we could consider the average over all sites. In a gapped system the best choice is the innermost site. In such a situation it might, however, be better to do a straight Lanczos calculation with $N_c + N_b$ cluster sites, instead of using N_b bath sites.

3.4 Symmetries

In the absence of spontaneous symmetry breaking the symmetries of the cluster (point-symmetries in CDMFT and additionally translation symmetry in DCA) are reflected in the Green matrix. In a symmetry-broken state with long-range order, like an antiferromagnet or a charge-density wave, the symmetry of the Green matrix is accordingly lowered. To exploit the symmetry we introduce vectors on the cluster that transform according to its irreducible representations. We write these vectors as $\mathbf{w}_{I,\nu}$ where I is the irreducible representation and $\nu = 1 \dots N_I$ counts linearly independent vectors transforming according to I . On an N_c -site cluster we can choose N_c such vectors that are orthonormal. Defining the matrix $\mathbf{W} = (\mathbf{w}_{I,\nu})$ of these vectors, we can block-diagonalize the bath Green matrix: $\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W}$ has blocks of dimension N_I corresponding to the different irreducible representations I . Since $\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W}$ is block diagonal for all ω , it follows from equation (57), that \mathbf{W} must also block-diagonalize the individual hybridization matrices $\mathbf{V}_l \mathbf{V}_l^\dagger$. Therefore the hybridization vectors must transform according to an irreducible representation: They can be written as $\mathbf{V}_l = \sum_\nu V_{l,I,\nu} \mathbf{w}_{I,\nu}$ for some irreducible representation I . If the \mathbf{V}_l also had components $\mathbf{w}_{J,\nu}$ of a different irreducible representation $J \neq I$ this would produce a hybridization matrix that could not be block-diagonalized. This can only happen for bath-sites with identical energy ε_l (accidental degeneracy): Assume \mathbf{V}_l and $\mathbf{V}_{l'}$ are the hybridizations for two bath-sites with $\varepsilon_l = \varepsilon_{l'}$. Then we can form arbitrary linear combinations of the hybridization matrices and hence of the hybridization vectors. For all these linear combinations the sum of the hybridization matrices must be block diagonal, and hence we can choose the hybridization vectors such that they transform according to irreducible representations.

We thus find that the bath-sites can be arranged into sets corresponding to the different irreducible representations. For fitting a block of the symmetrized bath Green matrix we need then only consider bath-sites of the respective irreducible representation. If the block is one-

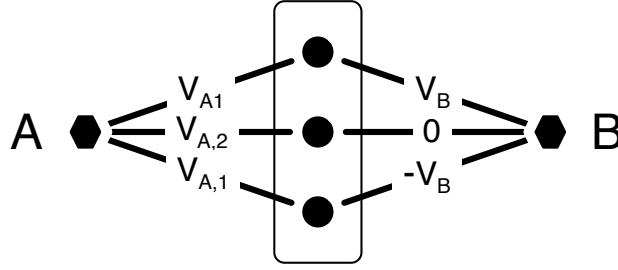


Fig. 13: Hybridization of bath-sites of symmetry A and B to a 3-site cluster. As defined in table 3, A is the unit representation, so a bath-site of type A has the same hybridization V to all cluster-sites that are equivalent by symmetry. B is the antisymmetric representation, so the hybridization of a bath-site of type B to cluster-sites that are related by mirror symmetry have the opposite sign. Consequently the hybridization to the central site of a linear cluster with an odd number of sites vanishes in the B representation.

dimensional we can choose the corresponding hybridizations real.

Sectors corresponding to different irreducible representations are only coupled through the Hubbard interaction U when solving the Anderson impurity model. Note that the coupling to bath-sites corresponding to an irreducible representation other than the unit representation lowers the symmetry of the impurity Hamiltonian with respect to that of the Green matrix.

CDMFT

As an example we consider a linear cluster of 3 sites as shown in figure 13. The symmetry is C_2 (see table 3). Transforming to the basis vectors $\mathbf{w}_{A,1} = (|1\rangle + |3\rangle)/\sqrt{2}$ and $\mathbf{w}_{A,2} = |2\rangle$ of symmetry A (see table 3) and $\mathbf{w}_B = (|1\rangle - |3\rangle)/\sqrt{2}$, we find the transformed bath Green matrix

$$\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W} = \begin{pmatrix} \mathcal{G}_{11}^{-1} + \mathcal{G}_{13}^{-1} & \sqrt{2} \mathcal{G}_{12}^{-1} & 0 \\ \sqrt{2} \mathcal{G}_{21}^{-1} & \mathcal{G}_{22}^{-1} & 0 \\ 0 & 0 & \mathcal{G}_{11}^{-1} - \mathcal{G}_{13}^{-1} \end{pmatrix}.$$

A bath-site of irreducible representation A contributes to the first block and has the same hybridization $V_{A,1}$ to the outer cluster-sites plus an independent hybridization parameter $V_{A,2}$ to the central site. A bath-site of irreducible representation B contributes to the second block. For such a bath-site the hybridization to cluster-sites that are related by mirror symmetry have opposite signs. Consequently, the hybridization to the central site vanishes.

The situation is slightly more complicated when the symmetry group has irreducible representations of dimension higher than one. The simplest example is the 2×2 cluster with C_{4v} symmetry. With $\mathbf{w}_{A_1} = (|1\rangle + |2\rangle + |3\rangle + |4\rangle)/2$, $\mathbf{w}_{B_2} = (|1\rangle - |2\rangle + |3\rangle - |4\rangle)/2$, and the pair $\mathbf{w}_{E,1} = (|1\rangle - |2\rangle - |3\rangle + |4\rangle)/2$, $\mathbf{w}_{E,2} = (|1\rangle + |2\rangle - |3\rangle - |4\rangle)/2$ we find that $\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W}$ is

C_2	E	σ_v	C_{2v}	E	C_2	σ_v	σ'_v
A	1	1	A_1	1	1	1	1
B	1	-1	A_2	1	1	-1	-1
			B_1	1	-1	1	-1
			B_2	1	-1	-1	1

C_{3v}	E	$2C_3$	$3\sigma_v$	C_{4v}	E	$2C_4$	C_4^2	$2\sigma_v$	$2\sigma_d$
A_1	1	1	1	A_1	1	1	1	1	1
A_2	1	1	-1	A_2	1	1	1	-1	-1
E	2	-1	0	B_1	1	-1	1	1	-1
				B_2	1	-1	1	-1	1
				E	2	0	-2	0	0

Table 3: Character tables of the point groups C_{1v} , C_{2v} , C_{3v} , and C_{4v} .

diagonal with diagonal elements

$$\begin{aligned}
(\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W})_{11} &= \mathcal{G}_{11}^{-1} + 2\mathcal{G}_{12}^{-1} + \mathcal{G}_{13}^{-1} \\
(\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W})_{22} &= \mathcal{G}_{11}^{-1} - 2\mathcal{G}_{12}^{-1} + \mathcal{G}_{13}^{-1} \\
(\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W})_{33} &= \mathcal{G}_{11}^{-1} - \mathcal{G}_{13}^{-1} \\
(\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W})_{44} &= \mathcal{G}_{11}^{-1} - \mathcal{G}_{13}^{-1}
\end{aligned}$$

A bath-site of symmetry A_1 has the same hybridization to all cluster-sites while for a bath-site of symmetry B_2 the hybridizations have alternating signs: $\mathbf{V}_l^\dagger = \bar{V}_l(1, -1, 1, -1)$. To realize the two-dimensional representation E we need two bath-sites l_1 and l_2 with degenerate energies $\varepsilon_{l_1} = \varepsilon_{l_2} = \varepsilon_l$ and hybridizations: $\mathbf{V}_{l_1}^\dagger = \bar{V}_l(1, -1, -1, 1)$ and $\mathbf{V}_{l_2}^\dagger = \bar{V}_l(1, 1, -1, -1)$. This is illustrated in figure 14.

DCA

As an example for DCA we consider a 3-site cluster with periodic boundary conditions. The symmetry group is C_{3v} (translations and inversion). Hence we introduce the basis vector $\mathbf{w}_{A_1} = (|1\rangle + |2\rangle + |3\rangle)/\sqrt{3}$, corresponding to $k = 0$, while the vectors formed by $\sin(2\pi/3)$ and $\cos(2\pi/3)$ give the E representation: $\mathbf{w}_{E,1} = (|1\rangle - |2\rangle)/\sqrt{2}$ and $\mathbf{w}_{E,2} = (|1\rangle + |2\rangle - 2|3\rangle)/\sqrt{6}$.

$$\mathbf{W}^\dagger \mathcal{G}^{-1} \mathbf{W} = \begin{pmatrix} \mathcal{G}_{11}^{-1} + 2\mathcal{G}_{12}^{-1} & 0 & 0 \\ 0 & \mathcal{G}_{11}^{-1} - \mathcal{G}_{12}^{-1} & 0 \\ 0 & 0 & \mathcal{G}_{11}^{-1} - \mathcal{G}_{12}^{-1} \end{pmatrix}.$$

In general bath-sites corresponding to the gamma point have the same hybridization to all cluster-sites, while those corresponding to $k = \pi$ have alternating hybridizations. For all other k -points we need two degenerate bath-sites, with hybridizations $V_{l_1,\mu} = V_l \sin(k\mu)$ and $V_{l_2,\mu} = V_l \cos(k\mu)$ to cluster-site μ .

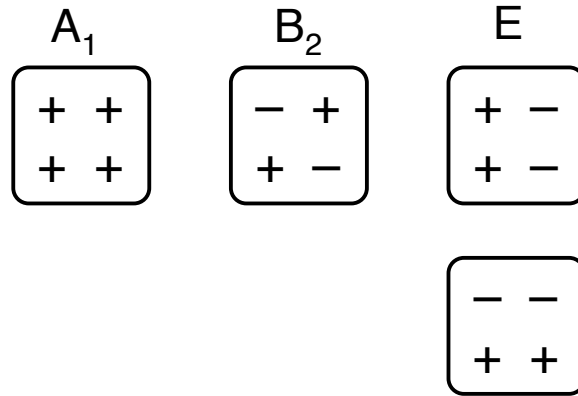


Fig. 14: Hybridization of bath-sites of symmetry A_1 , B_2 , and E to a 2×2 cluster. For a given irreducible representation the absolute value of the hybridization to all cluster-sites is the same, while the signs are indicated in the figure. Non-trivial hybridizations corresponding to irreducible representations A_2 or B_1 only appear for larger clusters.

4 Conclusions

We have seen that the Lanczos method is unbelievably efficient for calculating ground-state and dynamical response functions of many-body Hamiltonians. The determination of the ground-state takes only about $\mathcal{O}(\dim(\mathcal{H}))$ in time and memory. The iteration already converges after about a hundred steps, even for Hilbert spaces with dimensions in the trillions. This astounding convergence is based on the idea of steepest-descent to the ground state, which the Lanczos method even improves upon. In addition, we can very efficiently calculate Green functions. Here the rapid convergence is due to the fact that the Lanczos iteration reproduces more and more moments of the spectral function.

The great advantages of the Lanczos approach to strongly correlated systems is that it provides us with expressions for the Green function on the entire complex plane, i.e., in particular for real frequencies. Calculations are directly for the ground-state, i.e., $T=0$, but can easily be extended to finite, but low, temperatures.

The greatest disadvantage is the need to store full many-body vectors. Calculations are therefore restricted by the available memory to relatively small systems. To minimize the effects of finite system size it is therefore crucial to (i) efficiently use the vast distributed memories of current massively parallel machines and to (ii) find bath-parametrizations that minimize the effect of truncating it to finite size.

Acknowledgment

Support of the Deutsche Forschungsgemeinschaft through FOR1346 is gratefully acknowledged.

References

- [1] C. Lanczos: *An Iteration Method for the Solution of the Eigenvalue Problem of Linear Differential and Integral Operators*, J. Res. Nat. Bur. Stand. **49**, 255 (1950)
- [2] W.E. Arnoldi, Quarterly of Applied Mathematics **9**, 17 (1951)
- [3] C.C. Paige: *The Computation of Eigenvalues and Eigenvectors of Very Large Sparse Matrices* (PhD thesis, London University, 1971)
- [4] L.N. Trefethen and D. Bau III: *Numerical Linear Algebra* (Society for Industrial and Applied Mathematics, Philadelphia, 1997)
- [5] G.H. Golub and C.F. van Loan: *Matrix Computations* (Johns Hopkins University Press, 1996)
- [6] J.K. Cullum and R.A. Willboughby: *Lanczos Algorithms for Large Symmetric Eigenvalue Computations, Vol. 1: Theory, Vol. 2: Programs* (Birkhäuser, Boston, 1985)
- [7] R. Haydock: *The Recursive Solution of the Schrödinger Equation*, in H. Ehrenreich, F. Seitz, and D. Turnbull (eds.): *Solid State Physics* **35**, 216 (Academic Press, 1980)
- [8] E. Dagotto, Rev. Mod. Phys. **66**, 763 (1994)
- [9] L.V. Kantorovich, Uspekhi Mat. Nauk, **3**, 89 (1948)
- [10] M.R. Hestenes and W. Karush, J. Res. Nat. Bureau Standards **47**, 45 (1951)
- [11] A.N. Krylov, Izv. Akad. Nauk SSSR, Otd. Mat. Estest. **7**, 491 (1931)
- [12] J. Jaklič and P. Prelovšek, Adv. Phys. **49**, 1 (2000)
M. Aichhorn et al., Phys. Rev. B **67**, 161103 (2003)
- [13] A. Dolfen: *Massively parallel exact diagonalization of strongly correlated systems* (Diploma Thesis, RWTH Aachen University, 2006)
- [14] A. Dolfen, T.L. Luo, and E. Koch, Advances in Parallel Computing **15**, 601 (2008)
- [15] M. Caffarel and W. Krauth, Phys. Rev. Lett. **72**, 1545 (1994)
- [16] Q. Si, M.J. Rozenberg, G. Kotliar, and A.E. Ruckenstein, Phys. Rev. Lett. **72**, 2761 (1994)
- [17] M.J. Rozenberg, G. Moeller, and G. Kotliar, Mod. Phys. Lett. B **8**, 535 (1994)
- [18] G. Sangiovanni, A. Toschi, E. Koch, et al., Phys. Rev. B **73**, 205121 (2006)
- [19] E. Koch, G. Sangiovanni, and O. Gunnarsson, Phys. Rev. B **78**, 115102 (2008)
- [20] W. Kohn, Phys. Rev. **133**, A171 (1964)