

An efficient and accurate decomposition of the Fermi operator

Michele Ceriotti,^{a)} Thomas D. Kühne, and Michele Parrinello
*Computational Science, Department of Chemistry and Applied Biosciences, ETH Zurich, USI Campus,
 Via Giuseppe Buffi 13, CH-6900 Lugano, Switzerland*

(Received 25 March 2008; accepted 29 May 2008; published online 9 July 2008)

We present a method to compute the Fermi function of the Hamiltonian for a system of independent fermions based on an exact decomposition of the grand-canonical potential. This scheme does not rely on the localization of the orbitals and is insensitive to ill-conditioned Hamiltonians. It lends itself naturally to linear scaling as soon as the sparsity of the system's density matrix is exploited. By using a combination of polynomial expansion and Newton-like iterative techniques, an arbitrarily large number of terms can be employed in the expansion, overcoming some of the difficulties encountered in previous papers. Moreover, this hybrid approach allows us to obtain a very favorable scaling of the computational cost with increasing inverse temperature, which makes the method competitive with other Fermi operator expansion techniques. After performing an in-depth theoretical analysis of computational cost and accuracy, we test our approach on the density functional theory Hamiltonian for the metallic phase of the LiAl alloy. © 2008 American Institute of Physics. [DOI: 10.1063/1.2949515]

I. INTRODUCTION

When calculating the various ground state properties of fermionic systems, it is important to have fast and accurate ways of evaluating the density matrix. For noninteracting fermions, this amounts to calculating the Fermi function associated with the system's Hamiltonian \mathbf{H} . In many applications, \mathbf{H} is either of empirical nature or the result of a self-consistent density functional theory (DFT) calculation. The standard method for computing the density matrix requires diagonalizing \mathbf{H} , an operation whose computational complexity scales cubically with the number of electronic degrees of freedom N . Having a linear scaling scheme to obtain this quantity is a key step for modeling larger systems, thus making possible the computational study of a vast class of problems, whose behavior cannot be described by smaller models. Areas in which such a technique would have a major impact include nanotechnology and biochemistry, to name but a couple.

Several methods have been proposed to circumvent diagonalization.¹ These methods are based on the nearsightedness principle,^{2,3} which guarantees that in the $N \rightarrow \infty$ limit, the matrices needed to compute the Fermi operator will become sparse. Among the different approaches that have been proposed, we might mention divide-and-conquer,⁴ density-matrix minimization,⁵ and penalty methods,² as well as schemes based on Green's functions⁶ and maximally localized orbitals.⁷ The use of sparse matrix algebra eventually leads to linear scaling both in terms of memory requirements and of computational cost. A second class of methods, on which we shall focus here, uses the finite-temperature Fermi operator. Due to the finite temperature, the singularity at the chemical potential μ is smoothed, thus allowing for an expansion in simpler functions of \mathbf{H} . Since orbital localization

is not explicitly exploited, this class of methods can also be applied to metals. The earliest attempts in this direction were based on an expansion in Chebyshev polynomials.^{8,9} The computational cost of this method has been analyzed by Baer and Head-Gordon,¹⁰ who found that the order m of the polynomial needed to achieve a 10^{-D} accuracy depends linearly on the width of the Hamiltonian spectrum ΔE and the electronic temperature $1/\beta$, i.e., $m \sim D\beta\Delta E$. This obviously raises some problems when considering Hamiltonians with large ΔE , such as those arising from DFT calculations using plane wave basis sets or when low electronic temperatures are required. Recently it has been suggested¹¹ that fast polynomial summation methods, requiring a number of multiplications $\sim \sqrt{m}$, can be applied to the Fermi operator expansion, leading to the more favorable scaling $\sqrt{\beta\Delta E}$.

In this paper we revisit a particular form for the expansion of the Fermi operator, which is based on the grand-canonical formalism and developed in a series of recent papers.¹²⁻¹⁵ The grand-canonical potential for independent fermions is split into a sum of P terms, containing $e^{-\beta(\mathbf{H}-\mu)/2P}$. As a consequence of this decomposition, the Fermi operator can be written exactly as a sum of P terms. The larger the number of terms, the easier the evaluation of the exponential: this implies a trade-off between the size of P and the accuracy of the results. In this paper, we investigate the analytical properties of this decomposition, finding that a large number of terms are almost ideally conditioned, and that their contribution to the Fermi operator can be easily and effectively computed in a single shot with a polynomial expansion. The remaining few are tackled via a Newton-like iterative inversion scheme, which needs to be applied to each term individually but is very efficient in dealing with large $D\beta\Delta E$. With this hybrid approach, large values of P can be reached at a cost that is modest and independent of the system size. This result can improve significantly the prefactor of other methods using similar decompositions.¹²⁻¹⁵ More-

^{a)}Electronic mail: michele.ceriotti@phys.chem.ethz.ch.

over, using this approach, we achieve a scaling of the operation count with $D\beta\Delta E$ that is sublinear and competitive with the result of Ref. 11 if their fast summation technique is used. In this way, accurate, low-temperature calculations can be performed.

II. PROPERTIES OF THE EXPANSION

We use an expansion of the Fermi operator based on grand-canonical formalism, which has been developed and employed in several recent works.^{12–15} We summarize the derivation and the resulting expressions here, introducing a slightly different notation. To simplify, we will set the zero of energy at μ and measure energies in units of $k_B T$. This amounts to replacing in the standard expression for the Fermi operator $\beta(\mathbf{H}-\mu\mathbf{1})$ with \mathbf{H} . Using this notation, the grand-canonical potential for a system of noninteracting fermions becomes^{16,17}

$$\Omega = -2 \ln \det(\mathbf{1} + e^{-\mathbf{H}}) = -2 \text{Tr} \ln(\mathbf{1} + e^{-\mathbf{H}}). \quad (1)$$

Introducing the \mathbf{M}_l matrices,

$$\mathbf{M}_l = \mathbf{1} - e^{i(2l-1)\pi/2P} e^{-\mathbf{H}/2P}, \quad (2)$$

we can perform the decomposition

$$(\mathbf{1} + e^{-\mathbf{H}}) = \prod_{l=1}^P \mathbf{M}_l \mathbf{M}_l^*. \quad (3)$$

These expressions are analogous to those introduced in Ref. 12, apart from a change in indices ($P/2 \rightarrow P$, $P+1/2-l \rightarrow l$).

Using factorization (3), the grand-canonical potential can be written in compact form as $\Omega = -2 \text{Tr} \sum_{l=1}^P \ln(\mathbf{M}_l \mathbf{M}_l^*)$. The observables of interest for the system can be obtained as derivatives of the grand-canonical potential. In particular, the grand-canonical density matrix reads

$$\rho = \frac{\delta \Omega}{\delta \mathbf{H}} = \frac{1}{\mathbf{1} + e^{-\mathbf{H}}} = \frac{2}{P} \sum_{l=1}^P \mathbf{1} - \text{Re} \mathbf{M}_l^{-1}. \quad (4)$$

The decomposition (4) is exact for any value of P . As P increases, the exponential $e^{-\mathbf{H}/2P}$ is easier to approximate. However, the number of \mathbf{M}_l 's which have to be inverted increases. Previous works using this approach had to find the best compromise between the length of the expansion and the errors introduced by an approximate evaluation of the matrix exponential, therefore losing the advantage of an exact expansion. In order to find a solution to this problem, it is useful to analyze the properties of the \mathbf{M}_l 's in the large P limit. It turns out that matrices with small l are much more difficult to handle than those having a higher index. We therefore suggest applying different strategies in the two cases.

A. Properties of \mathbf{M}_l matrices

Let us define the spectral radius of a matrix \mathbf{A} as the maximum modulus of its eigenvalues, $\sigma(\mathbf{A}) = \max_i |a_i|$, and its condition number $\kappa(\mathbf{A}) = \sigma(\mathbf{A})\sigma(\mathbf{A}^{-1})$. We then introduce the shorthands $\Delta\epsilon = \sigma(\mathbf{H})$, which is a measure of the width of the Hamiltonian's spectrum, and $\delta\epsilon = 1/\sigma(\mathbf{H}^{-1})$, which is of the order of the band gap in insulators and tends to zero for




FIG. 1. (Color online) Plot of $\sqrt{1 + e^{-x/P} - 2e^{-x/2P} \cos \pi l/P}$, which is equal to $M_l(x)$ [Eq. (5)] within $\mathcal{O}(P^{-1})$. The dashed line corresponds to the locus of local minima.

metals. With this notation, the condition number of the Hamiltonian is $\kappa(\mathbf{H}) = \Delta\epsilon/\delta\epsilon$. In this section, we will obtain the corresponding quantities for the \mathbf{M}_l 's. In particular, we will show that $\kappa(\mathbf{M}_l)$ does not depend on P in the large P limit and demonstrate that the \mathbf{M}_l 's are always better conditioned than the Hamiltonian.

Therefore, we consider how the spectrum of the \mathbf{H} is mapped by the function

$$M_l(x) = |1 - e^{i\pi(2l-1)/2P} e^{-x/2P}|. \quad (5)$$

It is readily found that for any P and x , $M_l(x)$ is a monotonically decreasing function of l . For fixed l , $M_l(|x|) \leq M_l(-|x|)$ and the minimum value is $M_l(x_{\min}) = \sin \pi(2l-1)/2P$, which is reached for $x_{\min} = -2P \ln \cos \pi(2l-1)/2P$. From the plot of $M_l(x)$ (Fig. 1), it is apparent that the region which can lead to ill-conditioned matrices is the one with $l \ll P$ and $x \ll P$, where the spectrum of \mathbf{M}_l can contain eigenvalues close to zero. In this region, an upper bound to the maximum eigenvalue is given by $M_l(-\Delta\epsilon)$, and an estimate of the minimum eigenvalue within $\mathcal{O}(1/P)$ is $M_l(\delta\epsilon)$.

The following set of results can be easily proved by a series expansion in powers of $1/P$, assuming $l \ll P$ and $x \ll P$:

$$\sigma(\mathbf{M}_l) = \frac{1}{2P} \sqrt{\Delta\epsilon^2 + \pi^2(2l-1)^2} + \mathcal{O}(P^{-2}), \quad (6)$$

$$\sigma(\mathbf{M}_l^{-1}) = \frac{2P + \delta\epsilon/2}{\sqrt{\delta\epsilon^2 + \pi^2(2l-1)^2}} + \mathcal{O}(P^{-1}), \quad (7)$$

$$\begin{aligned} \kappa(\mathbf{M}_l) &= \sqrt{\frac{\Delta\epsilon^2 + \pi^2(2l-1)^2}{\delta\epsilon^2 + \pi^2(2l-1)^2}} + \mathcal{O}(P^{-1}) \\ &\sim \kappa(\mathbf{H}) \sqrt{\frac{\Delta\epsilon^2 + \pi^2(2l-1)^2}{\Delta\epsilon^2 + \kappa(\mathbf{H})^2 \pi^2(2l-1)^2}} \\ &\stackrel{\delta\epsilon \rightarrow 0}{\approx} 1 + \frac{\Delta\epsilon}{\pi(2l-1)}. \end{aligned} \quad (8)$$




FIG. 2. (Color online) Condition number of \mathbf{M}_l in the $P \rightarrow \infty$ limit for a typical value of $\Delta\epsilon$. Dark (blue) and light (red) series correspond to the behavior for a metal and for an insulator ($\delta\epsilon=20$). Even for a metallic system the condition number remains finite, and for the insulator, it saturates at $\kappa(\mathbf{H})$. In both cases, $\kappa(\mathbf{M}_l)$ drops rapidly to 1 as l increases.

It can be seen from Eq. (8) that the condition number $\kappa(\mathbf{M}_l)$ tends rapidly to 1 as l is increased and is always smaller than $\kappa(\mathbf{H})$ (see also Fig. 2). Note that the last inequality in Eq. (8), valid for $\delta\epsilon \rightarrow 0$, shows that $\kappa(\mathbf{M}_l)$ is bounded also in the metallic case.

III. A HYBRID APPROACH TO THE EXPANSION

The analysis performed above suggests dealing separately with the few, worst-conditioned \mathbf{M}_l matrices having $l < \bar{l}$ and with those which have $\kappa(\mathbf{M}_l) \sim 1$ for $l \geq \bar{l}$. The latter will form the “tail” contribution to the density matrix and will be discussed first.

A. Series expansion for the tail

In order to obtain a convergent power series for \mathbf{M}_l^{-1} , it is convenient to perform an expansion around the diagonal matrix $z'\mathbf{1}$, where z' is an arbitrary complex number whose value will be chosen so as to accelerate convergence. Defining the shorthand $\mathbf{Z}_l = e^{i\pi(2l-1)/2P} e^{-\mathbf{H}/2P}$, one has

$$\begin{aligned} \mathbf{M}_l^{-1} &= (1-z')^{-1} \left[\mathbf{1} - \frac{\mathbf{Z}_l - z'\mathbf{1}}{1-z'} \right]^{-1} \\ &= (1-z')^{-1} \sum_{j=0}^{\infty} \left(\frac{\mathbf{Z}_l - z'\mathbf{1}}{1-z'} \right)^j. \end{aligned} \quad (9)$$

The condition for convergence of Eq. (9) is that the whole spectrum of $(\mathbf{Z}_l - z'\mathbf{1})/(1-z')$ lies within the unit circle in the complex plane. Moreover, the convergence speed of the expansion will be determined by the eigenvalue which lies farthest from the origin (see Fig. 3). We refer to Appendix A for a detailed analysis of the convergence ratio:

$$\chi = \sigma \left(\frac{\mathbf{Z}_l - ke^{i\phi_l}\mathbf{1}}{1 - ke^{i\phi_l}} \right), \quad (10)$$

where we have set $z' = ke^{i\phi_l}$, defining $\phi_l = \pi(2l-1)/2P$ and introducing the complex-valued parameter k . There we show that in the large P limit, one obtains an upper bound to the convergence ratio, i.e., $\chi = \Delta\epsilon / \sqrt{\Delta\epsilon^2 + \pi^2(2l-1)^2} < 1$, provided one chooses for the optimal k the analytical estimate




FIG. 3. (Color online) The picture sketches the transformations in the complex plane leading from the Hamiltonian spectrum (1) to \mathbf{Z}_l (2), to $\mathbf{Z}_l - z'\mathbf{1}$ (3), and eventually to $(\mathbf{Z}_l - z'\mathbf{1})/(1-z')$ (4). The translation (2) \rightarrow (3) and the scaling (3) \rightarrow (4) depend both on the choice of z' . As described in the text, it is always possible to choose the parameter so as to keep the whole spectrum within the unit circle, ensuring convergence of the power series (9).

$$k = 1 - \frac{\Delta\epsilon^2}{2P\pi(2l-1)}. \quad (11)$$

Having ensured that the series (9) converges, we can estimate the error made by truncating the power series after m_T terms,

$$\begin{aligned} &\sigma \left((1-z')^{-1} \sum_{j=0}^{m_T} \left(\frac{\mathbf{Z}_l - z'\mathbf{1}}{1-z'} \right)^j - \mathbf{M}_l^{-1} \right) \\ &\leq \frac{1}{|1-z'|} \sum_{j=m_T+1}^{\infty} \sigma \left(\frac{\mathbf{Z}_l - z'\mathbf{1}}{1-z'} \right)^j = \frac{1}{|1-z'|} \frac{\chi^{m_T+1}}{1-\chi}. \end{aligned} \quad (12)$$

In order to achieve a 10^{-D} relative accuracy on \mathbf{M}_l^{-1} , it is necessary to retain at least

$$m_T \approx \frac{1}{\ln \chi} \left[\ln \left(\frac{1}{\chi} - 1 \right) - D \ln 10 + \ln(|1-z'| \sigma(\mathbf{M}_l^{-1})) \right]$$

terms. If we use Eqs. (11) and (7), setting $\delta\epsilon=0$, and taking the large $\Delta\epsilon$ limit, this estimate takes the simpler form

$$m_T \approx 2D \frac{\Delta\epsilon^2}{\pi^2(2l-1)^2} \ln 10. \quad (13)$$

While the scaling with $\Delta\epsilon^2$ is not optimal, the dependence on l^{-2} limits its effects to the small- l terms. These terms can be dealt with effectively with a different approach, as we will show below. The influence of the $\Delta\epsilon^2$ scaling on the overall operation count will therefore be limited.

Thanks to the chosen z' parametrization, the matrix powers entering Eq. (9) depend on l only by a scalar factor,

$$\frac{\mathbf{Z}_l - z'\mathbf{1}}{1-z'} = \frac{e^{i\phi_l}}{1 - ke^{i\phi_l}} (e^{-\mathbf{H}/2P} - k\mathbf{1}).$$

Therefore, we can compute the expensive powers $(e^{-\mathbf{H}/2P} - k\mathbf{1})^j$ just once and obtain any \mathbf{M}_l^{-1} by combining them with the appropriate scalar coefficients. Furthermore,




FIG. 4. (Color online) The number of terms required to achieve 10^{-3} relative accuracy in the polynomial expansion of \mathbf{M}_l^{-1} is plotted for a Hamiltonian with minimum eigenvalue -5 , with maximum eigenvalue 10 , and for $P=10^4$. A full line corresponds to results computed keeping k fixed to the $l=1$ value, while dots correspond to the results computed by optimizing k separately for each value of l . Dark (blue) and light (red) series correspond, respectively, to the results based on the analytical estimate (11) for k and to the ones obtained by iteratively minimizing Eq. (A1). Iterative refinement leads to a significant boost in performance. In any case, the number of terms computed for $l=1$ largely exceeds the terms needed to compute the contributions for larger l values, even if k is not optimized on a case-by-case basis.

we often need just the overall contribution to the density matrix arising from the tail, which reads

$$\mathbf{T}_{\bar{l}} = \sum_{l=\bar{l}}^P \mathbf{M}_l^{-1} = \sum_{j=0}^{m_T(\bar{l})} (e^{-H/2P} - k\mathbf{1})^j \sum_{l=\bar{l}}^P \frac{(1 - ke^{i\phi_l})^{-j}}{e^{-i\phi_l} - k}. \quad (14)$$

If either m_T or P is very large, computing the scalar coefficients in Eq. (14) implies a sizable overhead, which is, however, independent of the system size and becomes negligible for large systems.

In order to assess the accuracy of Eq. (14), further analysis is needed. If we want to reuse the powers $(e^{-H/2P} - k\mathbf{1})^j$, we must keep k fixed to the value optimized at \bar{l} . Expression (13) gives the number of terms required to compute \mathbf{M}_l^{-1} with 10^{-D} accuracy, *provided that k is optimized for each l* . However, the dependence of m_T on l offsets the effect of using a nonoptimal k . It is easy to show, given the estimate (11), that the number of terms computed for \bar{l} largely exceeds the number of terms required to compute \mathbf{M}_l^{-1} for any $l > \bar{l}$, even if k is kept fixed to the value optimized for \bar{l} . Figure 4 shows that this is the case also when k is iteratively optimized starting from the analytical estimate.¹¹

B. Newton inversion in the small- l region

To address the inversion of the worst-conditioned terms with $l < \bar{l}$, which are too expensive to obtain by polynomial expansion, one could resort to one of the techniques described in our previous work.^{12–15} In fact, the analysis performed so far can be seen as an improvement to those methods, since we can evaluate in one shot the contribution from the tail, lowering the number of terms which must be treated individually, and therefore improving the efficiency.

In this section we will discuss an alternative approach for computing the small- l \mathbf{M}_l^{-1} , based on a well-established Newton method for matrix inversion. We give a brief outline

of the algorithm and some of its known analytical properties,¹⁸ and will use them to estimate the number of operations necessary for our purposes. Given a nonsingular, $M \times M$ matrix \mathbf{A} , the iterative procedure

$$\mathbf{B}_{k+1} = 2\mathbf{B}_k - \mathbf{B}_k \mathbf{A} \mathbf{B}_k \quad (15)$$

converges to \mathbf{A}^{-1} . Defining $\mathbf{R}(\mathbf{B}) = \mathbf{1} - \mathbf{B}\mathbf{A}$, the condition for convergence is that $\chi = \sigma(\mathbf{R}(\mathbf{B}_0)) < 1$, and the error after k iterations is

$$\sigma(\mathbf{A}^{-1} - \mathbf{B}_k) = \sigma(\mathbf{B}_0) \chi^{(2^k)} (1 - \chi)^{-1}, \quad (16)$$

which corresponds to a number of multiplies (two per iteration)

$$m_N = \frac{2}{\ln 2} \ln \frac{\ln[10^{-D}(1 - \chi)\sigma(\mathbf{A}^{-1})/\sigma(\mathbf{B}_0)]}{\ln \chi} \quad (17)$$

needed to achieve a 10^{-D} relative accuracy.

One must then face the problem of finding the approximate inverse \mathbf{B}_0 needed to start the iterations (15). The authors of Ref. 18 suggested the simple form

$$\mathbf{B}_0 = \mathbf{A}^\dagger (\|\mathbf{A}\|_1 \|\mathbf{A}\|_\infty)^{-1}, \quad (18)$$

where $\|\mathbf{A}\|_1 = \max_j \sum_{i=1}^M |A_{ij}|$ and $\|\mathbf{A}\|_\infty = \max_i \sum_{j=1}^M |A_{ij}|$. If one uses Eq. (18), convergence is guaranteed. Taking as usual the large P and $\Delta\epsilon$ limit for a metallic system, one obtains $m_N \sim \ln \Delta\epsilon + \ln(D \ln 10) + \ln M / \pi^2 (2l-1)^2$ as an estimate of the operation count to invert \mathbf{M}_q . Even if a feeble M dependence has been introduced in the operation count, the efficiency is greatly improved if one needs high accuracy or if $\Delta\epsilon$ is large, thanks to the exponential convergence rate.

It is, however, more effective to exploit the simple analytic form for \mathbf{M}_l^{-1} to construct better initial guesses. For instance, one can use the following relation between \mathbf{M}_l^{-1} and $\mathbf{M}_{l-\delta l}^{-1}$:

$$\begin{aligned} \mathbf{M}_{l-\delta l}^{-1} &= \mathbf{M}_l^{-1} e^{i\pi\delta l/P} [\mathbf{1} + \mathbf{M}_l^{-1} (e^{i\pi\delta l/P} - \mathbf{1})]^{-1} \\ &= e^{i\pi\delta l/P} \sum_{j=0}^{\infty} (e^{i\pi\delta l/P} - \mathbf{1})^j \mathbf{M}_l^{-1(j+1)}, \end{aligned} \quad (19)$$

to get an estimate for $\mathbf{M}_{l-\delta l}^{-1}$ starting from an already-computed inverse. The series (19) converges provided that $|e^{i\pi\delta l/P} - 1| \sigma(\mathbf{M}_l^{-1}) < 1$. In the $P \rightarrow \infty$ limit, this amounts to the condition $\delta l < l - 1/2$. In theory, all the terms up to $l=1$ could be computed by inserting any \mathbf{M}_l^{-1} into Eq. (19). In practice, computing powers of \mathbf{M}_l^{-1} is not advisable if we aim at linear scaling since the \mathbf{M}_l^{-1} and their powers tend to be much fuller than the Hamiltonian, and the asymptotic convergence rate of Eq. (19) is worse than the one for the iterative inversion. In any case, the lowest-order approximation is already much more effective than the universal guess described in Ref. 18. One finds that the convergence ratio for the computation of \mathbf{M}_{l-1}^{-1} using the low-order extrapolation $e^{i\pi/P} \mathbf{M}_l^{-1}$ is $\chi \sim 2\pi / \sqrt{\delta\epsilon^2 + \pi^2(2l-1)^2}$, leading to an estimate for the operation count:




FIG. 5. (Color online) Total number of matrix-matrix multiplications required to obtain the density matrix, combining series expansion and Newton inversion methods, on a log-log plot. Light (red) and dark (blue) lines correspond to 10^{-5} and 10^{-8} target accuracies, respectively. Full (a), dashed (b), and dotted (c) lines correspond, respectively, to the number of operations estimated using the general-purpose expression (18), using a zeroth-order extrapolation guess, and using extrapolation together with fast polynomial evaluation in the tail region. Grid lines mark the slope expected for a linear dependence between $\Delta\epsilon$ in units of $k_B T$ and the overall operations count m_{tot} .

$$m_N = \frac{2}{\ln 2} \ln \frac{D \ln 10 - 2 \ln \pi}{\ln q} + \mathcal{O}(1/q). \quad (20)$$

This estimate is independent of $\delta\epsilon$ because we considered the worst-case scenario where the system is metallic. It is also independent of M and, most importantly, of $\Delta\epsilon$. In practice, one starts from $\mathbf{M}_{\bar{l}}^{-1}$ obtained from the polynomial expansion, then computes $\mathbf{M}_{\bar{l}-1}^{-1}$ using $e^{i\pi/P}\mathbf{M}_{\bar{l}}^{-1}$ as the initial guess, and continues stepwise, obtaining the initial estimate for iterative inversion of $\mathbf{M}_{\bar{l}-2}$ from the previously computed $\mathbf{M}_{\bar{l}-1}^{-1}$, and so on. Alternatively, the first inverse matrix $\mathbf{M}_{\bar{l}-1}^{-1}$ can be computed starting from the simple guess Eq. (18). Efficient higher-order extrapolations will be discussed in Appendix B.

C. Overall operation count

In the previous section we obtained [Eqs. (13) and (17)] an upper bound estimate of the number of matrix-matrix multiplications needed in order to obtain the tail contribution up to \bar{l} and to invert a single \mathbf{M}_l using an iterative Newton method. The optimal value for \bar{l} is obtained when the incremental cost of including an extra term in the tail contribution \mathbf{T}_l [cf. Eq. (14)] becomes larger than the cost of a single iterative inversion, i.e., when

$$m_T(\bar{l}) - m_T(\bar{l} + 1) \geq m_N(\bar{l}). \quad (21)$$

The overall number of multiplications is then

$$m_{\text{tot}} = m_T(\bar{l}) + \sum_{l=1}^{\bar{l}-1} m_N(l). \quad (22)$$

In Fig. 5 we plot the overall operation count obtained by using our theoretical estimates for m_T and m_N . A dramatic improvement is obtained when we use $e^{i\pi/P}\mathbf{M}_{\bar{l}}^{-1}$ as the initial guess for the inversion of $\mathbf{M}_{\bar{l}-1}$. We can think of the extrapo-

lated guess as an almost optimal preconditioner and are considering how this could be exploited in different inversion schemes as well. It is worth noting that despite the fact that the tail contribution requires a number of multiplies scaling quadratically with $\Delta\epsilon$, the overall scaling is significantly sub-linear. Comparing our results [Fig. 5(b)] with the multiplication count for standard Chebyshev polynomial expansion, as given by Ref. 10, our method becomes beneficial at $\Delta\epsilon \sim 20$ —the break-even point getting lower as the target accuracy D is increased. Fast polynomial summation methods^{11,19,20} can be used to compute both \mathbf{M}_l^{-1} and \mathbf{T}_l . This reduces the number of multiplies from m_T to $3\sqrt{m_T}$, however, at the cost of storing extra $\sqrt{m_T}$ matrices. Combining these fast summation techniques with iterative inversion further lowers the operation count, leading to a scaling slightly better than $\sqrt{\Delta\epsilon}$ [Fig. 5(c)]. In this case, however, the prefactor of our method is larger, so that the break-even point, when comparing with Refs. 11, 19, and 20, is shifted towards higher accuracy and larger $\Delta\epsilon$. We are currently investigating the possibility of applying an alternative expansion of the tail contribution, which should make both our $\Delta\epsilon$ scaling and the prefactor highly competitive.

IV. A TEST CASE

So far we have estimated the accuracy of the computation of each \mathbf{M}_l^{-1} term using $\epsilon(\tilde{\mathbf{M}}_l^{-1}) = \sigma(\mathbf{M}_l^{-1} - \tilde{\mathbf{M}}_l^{-1}) / \sigma(\mathbf{M}_l^{-1})$ as a measure of the error affecting the estimate $\tilde{\mathbf{M}}_l^{-1}$. However, the quantity we are more interested in is the band-structure energy $E = \text{Tr}[\rho\mathbf{H}]$. A theoretical estimation of the error on E requires several assumptions on the distribution of errors over the different eigenvalues of the Hamiltonian and the different l terms, and we have not attempted it here. We have instead tested our method against a real system, selecting the self-consistent DFT Hamiltonian matrix of a 128-atom sample of the metallic fcc phase of LiAl, as computed by the CP2K (Refs. 21 and 22) package.²³ The orthogonal Hamiltonian matrix is obtained by multiplying the nonorthogonal one with the inverse square root of the overlap matrix.²⁴ We then computed with standard diagonalization techniques the chemical potential and the exact band-structure energy for different electronic temperatures. We also obtained the bounds of the spectrum of $\mathbf{H}(\epsilon_+ = 121.15 \text{ eV}$ and $\epsilon_- = -42.65 \text{ eV})$, which are needed in Eq. (A1) and could in principle be computed in linear scaling with the Lanczos method or easily estimated by Gershgorin's circle theorem²⁵ or any matrix norm.

We then applied our algorithm to the orthogonalized Hamiltonian using fast polynomial summation to compute the tail and using first-order extrapolation in the Newton region, with a history vector containing the last two matrices [cf. Eq. (B1)]. Slight improvements in the operation count could be obtained by hand tuning \bar{l} , but we just used the automatic procedure based on our theoretical estimates, as described in the previous section. In Fig. 6 we plot the number of multiplications performed versus the resulting error on the energy. Since we can use a large value of P , $e^{-H/2P}$ can be computed with only a few matrix-matrix multiplies, which have not been included in the operation count.




FIG. 6. (Color online) Number of matrix-matrix multiplications used to achieve a given error on the band-structure energy for different electronic temperatures. Details of the system are given in the text. The data points for every temperature, from left to right, correspond to 10^{-2} , 10^{-3} , 10^{-5} , and 10^{-7} target accuracies.

For a given target accuracy, the operation count scales better than $\sqrt{\beta}$ (Fig. 7). We also observe that the accuracy of the energy is much better than the relative accuracy guaranteed by the theoretical estimates. Consider, for example, that, by requiring a relative “spectral radius accuracy” better than 10^{-2} (first data points in Fig. 6), we obtain a relative error on the energy of the order of 10^{-4} (the total energy is ~ -5 keV). This is mainly due to the fact that the error in the energy is second order with respect to the error in the density matrix. However, we observe that also the error in the full density matrix, computed as the spectral radius of the difference with the result obtained with diagonalization, is in general almost one order of magnitude smaller than the required accuracy. This result is probably due to a combination of effects: firstly, we use worst-case estimates, so that the accuracy of the individual terms is necessarily higher than the assumed one. Moreover, the errors affecting different l terms might partially cancel out each other, and many of the contributions in the Newton region are computed with an accuracy much higher than requested due to the exponential convergence. The accuracy improves very quickly as the number of operations increases until, for errors around




FIG. 7. (Color online) The number of matrix-matrix multiplies performed to compute the density matrix for the LiAl test case for different electronic temperatures and target accuracies, plotted on a log-log scale, together with guidelines corresponding to a $m \propto \beta \Delta \epsilon$ scaling. Here $\Delta \epsilon = 82$ eV.

0.01 meV/atom, numerical issues come into play and prevent further refinement, which is anyway hardly necessary for most applications.

Most of the observables relevant to electronic structure calculations, such as forces and electronic density, are readily evaluated by expressions of the form $\langle A \rangle = \text{Tr}[\rho \mathbf{A}]$. Since the matrix \mathbf{A} obeys the same sparsity as the Hamiltonian, $\langle A \rangle$ depends only on a small subset of the nonzero elements of the density matrix. We are currently investigating whether it is possible to compute the expectation value directly, without evaluating nonrelevant elements of ρ , which would further improve the efficiency.

V. CONCLUSION

We have performed a detailed study of a recently proposed form for the Fermi operator expansion. The properties of this expansion allow features of the expansion in polynomial and rational functions to be combined and by optimizing the mixture we can have the best of both worlds. In this way, we circumvent the trade-off between the number of terms and the accuracy of the expansion, which was needed by prior implementations of this expansion of the Fermi operator. Moreover, sublinear scaling of the matrix-matrix multiplication count with respect to the Hamiltonian range is achieved, making the method particularly attractive for low-temperature and high-accuracy applications. However, there is still room for improvement. In particular, work is in progress in the direction of a better polynomial expansion in the tail region. We are also considering applying the method to *ab initio* molecular dynamics. In this case one could use the \mathbf{M}_l^{-1} 's stored from the previous steps as a starting point for iterative minimization.³⁰ In this way, the computation of the different l channels can be made independent, adding a layer of parallelism on top of the parallel matrix-matrix multiply. Formal analogies between our expansion and Trotter factorization entering path integral techniques suggest that some of the ideas presented here might be useful to tackle that problem as well. In order to achieve linear scaling, attention should be paid to the issue of matrix truncation, since here we have dealt only with matrix-matrix operation counts. Preliminary results show that in this respect, there are no significant differences from standard expansion methods, as the minimum sparsity of the terms taken into account is basically the same as the sparsity of the whole density matrix, which is dictated by the physics of the system. In the end, we have obtained significant improvements over previous implementations of the Fermi operator decomposition, and laid solid analytical foundations for further progress.

ACKNOWLEDGMENTS

The generous allocation of computer time by the Swiss National Supercomputing Center (CSCS) and technical assistance from Neil Stringfellow is kindly acknowledged. We would also like to thank Giovanni Bussi, Paolo Elvati, and Luca Ferraro for fruitful discussion.

APPENDIX A: OPTIMAL PARAMETER FOR SERIES EXPANSION

We show how the value of k in Eq. (10) can be optimized to obtain faster convergence of the polynomial expansion. The expressions involved are quite lengthy, so we introduce several shorthands. Let ϵ_{\pm} be the bounds of the Hamiltonian spectrum. We parameterize k as $k=(1+e^{i\theta}r/P)$, define $e^{i\pi(2l-1)/2P}=(v_l+iw_l)$, and $s_{\pm}=e^{-\epsilon_{\pm}/2P}$. The square modulus of the extrema of the transformed Hamiltonian spectrum (see Fig. 3) is

$$d_{\pm}^2 = \frac{r^2 + P^2(s_{\pm} - 1)^2 - 2Pr(s_{\pm} - 1)\cos\theta}{r^2 + 2P^2(1 - v_l) + 2Pr(\cos\theta(1 - v_l) + w_l \sin\theta)} \quad (\text{A1})$$

and the convergence ratio is $\chi=\max(d_+, d_-)$. One can obtain an analytical estimate for k and an upper bound for χ by taking the $P \rightarrow \infty$ limit and making the simplifying assumption $|\epsilon_{\pm}|=\Delta\epsilon$. This implies $\theta=-\pi/2$ and leads to the estimate (11), which can be further improved by minimizing numerically Eq. (A1) with respect to θ and r .

APPENDIX B: HIGH-ORDER INITIAL GUESS FOR ITERATIVE INVERSION

One can derive expressions for high-order extrapolation of inverse \mathbf{M}_l matrices from Eq. (19), writing them as a linear combination of already-computed inverses. We will sketch the procedure by deriving the expression for the first-order extrapolation of \mathbf{M}_{l-1}^{-1} using only \mathbf{M}_l^{-1} and \mathbf{M}_{l+1}^{-1} , which is then easily extended to higher orders. Let $c_j^{(n)} = e^{-ni\pi/P}(e^{-ni\pi/P} - 1)^{j-1}$. One can write the first-order extrapolations for the new inverse and for the already-computed one as a function of powers of \mathbf{M}_l^{-1} :

$$\mathbf{M}_{l-1}^{-1} = c_1^{(-1)}\mathbf{M}_l^{-1} + c_2^{(-1)}\mathbf{M}_l^{-2},$$

$$\mathbf{M}_{l+1}^{-1} = c_1^{(1)}\mathbf{M}_l^{-1} + c_2^{(1)}\mathbf{M}_l^{-2}.$$

This linear system can be solved for \mathbf{M}_{l-1}^{-1} and \mathbf{M}_l^{-2} , obtaining

$$\mathbf{M}_{l-1}^{-1} \sim \mathbf{M}_l^{-1}(e^{i\pi/P} + e^{2i\pi/P}) - \mathbf{M}_{l+1}^{-1}e^{3i\pi/P}. \quad (\text{B1})$$

For higher orders one simply inserts into the system more constraints, corresponding to “older” inverse matrices, and writes the extrapolation including higher powers of \mathbf{M}_l^{-1} . The

system is then solved in terms of these powers, eventually finding the coefficients for the estimate of the new inverse as a linear combination of the older ones.

- ¹S. Goedecker, *Rev. Mod. Phys.* **71**, 1085 (1999).
- ²W. Kohn, *Phys. Rev. Lett.* **76**, 3168 (1996).
- ³E. Prodan and W. Kohn, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 11635 (2005).
- ⁴W. Yang, *Phys. Rev. Lett.* **66**, 1438 (1991).
- ⁵X. Li, R. Nunes, and D. Vanderbilt, *Phys. Rev. B* **47**, 10891 (1993).
- ⁶S. Baroni and P. Giannozzi, *Europhys. Lett.* **17**, 547 (1992).
- ⁷G. Galli and M. Parrinello, *Phys. Rev. Lett.* **69**, 3547 (1992).
- ⁸S. Goedecker and L. Colombo, *Phys. Rev. Lett.* **73**, 122 (1994).
- ⁹S. Goedecker and M. Teter, *Phys. Rev. B* **51**, 9455 (1995).
- ¹⁰R. Baer and M. Head-Gordon, *J. Chem. Phys.* **107**, 10003 (1997).
- ¹¹W. Liang, C. Saravanan, Y. Shao, R. Baer, A. Bell, and M. Head-Gordon, *J. Chem. Phys.* **119**, 4117 (2003).
- ¹²F. R. Krajewski and M. Parrinello, *Phys. Rev. B* **71**, 233105 (2005).
- ¹³F. R. Krajewski and M. Parrinello, *Phys. Rev. B* **73**, 041105 (2006).
- ¹⁴F. R. Krajewski and M. Parrinello, *Phys. Rev. B* **74**, 125107 (2006).
- ¹⁵F. Krajewski and M. Parrinello, *Phys. Rev. B* **75**, 235108 (2007).
- ¹⁶A. Alavi and D. Frenkel, *J. Chem. Phys.* **97**, 9249 (1992).
- ¹⁷A. Alavi, J. Kohanoff, M. Parrinello, and D. Frenkel, *Phys. Rev. Lett.* **73**, 2599 (1994).
- ¹⁸V. Pan and J. Reif, in *STOC'85: Proceedings of the 17th Annual ACM Symposium on Theory of Computing* (ACM, New York, NY, 1985), p. 143.
- ¹⁹W. Liang, R. Baer, C. Saravanan, Y. Shao, A. Bell, and M. Head-Gordon, *J. Comput. Phys.* **194**, 575 (2004).
- ²⁰C. Van Loan, *IEEE Trans. Autom. Control* **24**, 320 (1979).
- ²¹J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, *Comput. Phys. Commun.* **167**, 103 (2005).
- ²²See <http://cp2k.berlios.de> for source code and documentation.
- ²³We have used GTH pseudopotentials (Refs. 26–28), gradient-corrected exchange-correlation functional [PBE (Ref. 29)], and a double- ζ basis with one additional set of polarization functions, for a total of 1728 basis functions. LiAl in the fcc phase is a metal. Since we are computing the Hamiltonian at the Γ point only, the spectrum has six half-occupied degenerate states at the zero-temperature Fermi energy. In the low-temperature limit, $\delta\epsilon \rightarrow 0$, which makes this system particularly challenging.
- ²⁴P. Löwdin, *J. Chem. Phys.* **18**, 365 (1950).
- ²⁵A. Palsler and D. Manolopoulos, *Phys. Rev. B* **58**, 12704 (1998).
- ²⁶S. Goedecker, M. Teter, and J. Hutter, *Phys. Rev. B* **54**, 1703 (1996).
- ²⁷M. Krack, *Theor. Chem. Acc.* **114**, 145 (2005).
- ²⁸C. Hartwigsen, S. Goedecker, and J. Hutter, *Phys. Rev. B* **58**, 3641 (1998).
- ²⁹J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- ³⁰T. D. Kühne, M. Krack, F.R. Mohamed, and M. Parrinello, *Phys. Rev. Lett.* **98**, 066401 (2007).