An efficient GPU algorithm for tetrahedron-based Brillouin-zone integration

Daniel Guterding[1,*] and Harald O. Jeschke[2]

[1] *Lucht Probst Associates, Große Gallusstraße 9, 60311 Frankfurt am Main, Germany, European Union*
[2] *Research Institute for Interdisciplinary Science, Okayama University, Okayama 700-8530, Japan*

We report an efficient algorithm for calculating momentum-space integrals in solid state systems on modern graphics processing units (GPUs). We extend the tetrahedron method by Blöchl *et al.* to the more general case of the integration of a momentum as well as energy dependent quantity and implement the algorithm based on the CUDA programming framework. We test this method by applying it to a simple example, the calculation of the orbital-resolved density of states. We benchmark our code on the problem of calculating the orbital-resolved density of states in an iron-based superconductor and discuss the design choices made in the implementation. Our algorithm delivers large speedups of up to a factor ~ 165 also for moderately sized workloads compared to standard algorithms executed on central processing units (CPUs).

I. INTRODUCTION

When calculating the properties of periodic solid state systems, the Brillouin zone naturally appears as the smallest possible unit of momentum space that respects all spatial symmetries of the system. Many properties, like the total energy, the density of states or even magnetic susceptibilities involve integrations over the Brillouin zone. Consequently, the performance of the integration algorithm is crucial to calculating these properties efficiently.

A particularly appealing algorithm for Brillouin zone integrations employs linear approximations within tetrahedra constructed from a rectangular grid of data points in momentum space. Various tetrahedron methods have been developed to obtain the density of states [1, 2], magnetic susceptibility [3, 4], phonon response [5] and electron-phonon interactions [6]. Recently, advances for further many-body quantities have been reported [7].

In this paper we do not aim to qualitatively improve the calculation of one of these quantities, nor do we want to add a further quantity to the list of those that can be calculated using a tetrahedron method. Instead, we present a technique to accelerate tetrahedron integration using modern graphics processing units (GPUs). We show that tetrahedron algorithms are by construction memory-bound and, therefore, well suited for GPU implementation. Indeed, we achive speedups of up to a factor ~ 165, thereby vastly extending the range of accessible problem sizes.

As an example, we apply our algorithm to the problem of calculating the orbital-resolved density of states in multi-orbital systems. However, our algorithm can be applied to other quantities, such as the calculation of susceptibilities, with minimal modification. Therefore, we make our implementation in the CUDA programming framework available as an open-source code. We would like to encourage the scientific community to use our code and report any further opportunities for optimization.

We start our paper by introducing the model system that serves as an input to our algorithm. We then explain the general tetrahedron method and present our extension to integrating general momentum dependent quantities and, specifically, the orbital-resolved density of states. Subsequently, we outline our goals in designing an implementation of the tetrahedron method for GPUs and explain how we meet those goals in our specific implementation. We benchmark our implementation in two realistic scenarios and relate the benchmark results to the decisions made in our implementation. Finally, we summarize our findings.

II. MODELS AND METHODS

A. Model Hamiltonians

In this paper we present a method for calculating the orbital-resolved density of states in real materials. A class of compounds that has a density of states with reasonably complex orbital structure, are iron-based superconductors [8, 9]. In these materials, the tetragonal crystal field leads to a splitting of the Fe $3d$ orbitals into four distinct groups, which determine the physics at the Fermi level. Therefore, we use these iron-based materials as an example, where the orbital-structure of the density of states is of interest.

The material of choice is FeSe due to its simple crystal structure and the availability of accurate tight-binding models. Similar to Ref. [10] we construct an eight-orbital tight-binding model including Fe $3d$ and Se $4p$ states. The calculated hopping parameters are denoted as t_{ij}^{sp}, where i and j are lattice site indices and indices s and p identify the orbitals.

$$H_0 = -\sum_{i,j,s,p,\sigma} t_{ij}^{sp} c_{is\sigma}^{\dagger} c_{jp\sigma} \tag{1}$$

This Hamiltonian can be diagonalized as a function of momentum $\vec{k}$ by inserting the Fourier-transform of

* daniel.guterding@gmail.com; The program code is available at https://github.com/danielguterding/cutetra.

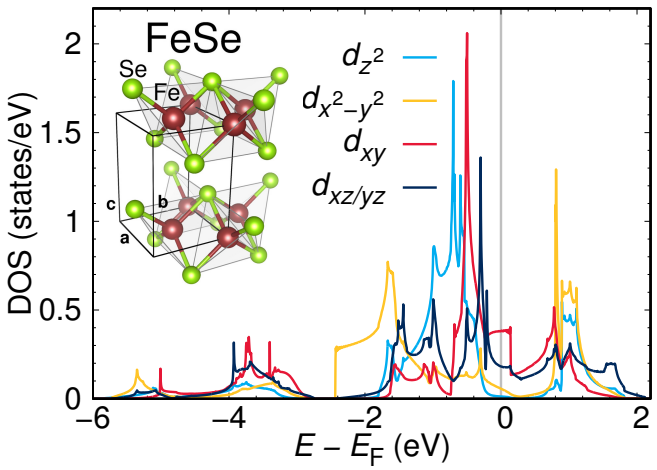

FIG. 1. (Color online) Orbital-resolved density of states of the eight-band model for FeSe. Shown is only the density of states for the Fe 3*d* orbitals. The inset shows the crystal structure of FeSe.

the operators and using standard matrix diagonalization techniques. This yields band energies $E_n(\vec{k})$ and matrix elements $a_n^m(\vec{k})$ that connect orbital and band space denoted by indices m and n respectively.

We are now interested in orbital dependent integrals over the Brillouin zone of the type

$$\langle X_m \rangle = \frac{1}{V_G} \sum_n \int_{V_G} d^3k \, X_n^m(\vec{k}) f\big(E_n(\vec{k})\big), \qquad (2)$$

where the occupation numbers $f(E)$ are given at $T = 0$ by the Heaviside step function $\Theta(E)$, and V_G is the volume of the Brillouin zone. The orbital-resolved particle number $n_m(E)$ would for example be obtained for the choice

$$X_n^m(\vec{k}) = b_{mn}(\vec{k}) \equiv a_n^{m*}(\vec{k}) a_n^m(\vec{k}). \qquad (3)$$

Then, an energy derivative leads to the orbital-resolved density of states

$$\rho_m(E) = \frac{1}{V_G} \sum_n \int_{V_G} d^3k \, b_{mn}(\vec{k}) \delta\big(E - E_n(\vec{k})\big). \qquad (4)$$

However, we have not specified yet how to calculate the momentum integrals in practice. Our method of choice is the so-called *tetrahedron method*, which we explain in the following. An example for the density of states of FeSe, calculated using our method, is shown in Fig. 1.

B. Generic tetrahedron method

The general idea of the tetrahedron method [1] is that a quantity $\langle X \rangle$, integrated over the crystal, can be obtained as a summation over discrete k-points with weights $w_{n,j}$ determined by a linear interpolation over the edges of the tetrahedra.

$$\langle X \rangle = \sum_n \int_{V_G} d\vec{k} \, X_n(\vec{k}) = \sum_{n,j} X_n(\vec{k}_j) w_{n,j} \qquad (5)$$

Here, j is the index of the k-point and n is the band index.

The formula can be rewritten in terms of tetrahedra, labelled by index i, which each contain four k-points, i.e.

$$\langle X \rangle = \sum_{i,n} T_{i,n}, \qquad (6)$$

where the tetrahedron contribution $T_{i,n}$ is given by

$$T_{i,n} = \sum_{l=1}^{4} X_n(\vec{k}_l) w_{n,l} \qquad (7)$$

and l denotes the k-point index within the tetrahedron.

Determination of an orbital dependent integral $\langle X_m \rangle$ now requires a small generalization of this standard approach as detailed below.

C. Tetrahedron integration of the orbital-resolved density of states

In the tetrahedron method for the total density of states the quantity to integrate X_n^m is set to $X_n^m = 1$ to simply count the number of states. Subsequently, the derivative with respect to the energy of interest E is taken to obtain the density of states. Since X_n^m is constant within the tetrahedron, the total contribution of the tetrahedron can be written down as a single expression as given in Ref. [1].

In our case, we want to integrate the orbital-resolved density of states. Therefore, the quantity to integrate X_n is now equal to the orbital weight $b_{n,l,m} \in [0, 1]$, which does depend on momentum, even within a tetrahedron. The orbital weight $b_{n,l,m}$ is given by the square of the matrix elements $a_n^m(\vec{k}_l)$ as given in Eq. 3. Since the the orbital weights $b_{n,l,m}$ are not constant within a tetrahedron (they depend on momentum indexed by l), the contributions from the four corner points within each tetrahedron need to be summed up explicitly, contrary to the case of the total density of states. For this reason, we need to derive the integration weights $w_{n,l}$ at each of the four corner points instead of just one weight w_n for the entire tetrahedron.

For the energy dependence, note that the integration weights $w_{n,l}$ depend on the energy of interest E, while the orbital weights $b_{n,l,m}$ only depend on the fixed energy and momentum indices n, l. Consequently, taking the derivative with respect to the energy of interest E in order to go from the number of states to the density of states amounts to taking the derivative of the integration weights $w_{n,l}$ only. Therefore, the tetrahedron contribution to the orbital-resolved density of states can be

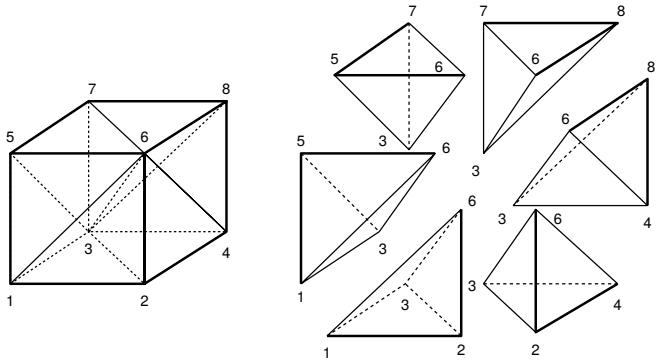

FIG. 2. (Color online) The left-hand side of the figure shows an elemental cube of the input k-grid with its corner points labelled from 1 to 8. The right-hand side of the figure shows how the cube is divided into six tetrahedra. The labels correspond to the corner labels on the left-hand side.

written as

$$T_{i,n,m}(E) = \sum_{l=1}^{4} b_{n,l,m} \frac{\partial w_{n,l}(E)}{\partial E}, \tag{8}$$

where m denotes the orbital. The orbital-resolved density of states $\rho_m(E)$ is then given as the sum over all tetrahedral contributions

$$\langle \rho_m(E) \rangle = \sum_{i,n} T_{i,n,m}(E). \tag{9}$$

The definitions of the tetrahedra are the same as in the Ref. [1]. For completeness, in Fig. 2 we show how an elemental cube on the input k-grid is divided into six tetrahedra. The weights are calculated from linear interpolation along the edges of the tetrahedra. The band energies at constant band index n at each of the four corners of a tetrahedron are sorted in ascending order and assigned labels $E_1 \leq E_2 \leq E_3 \leq E_4$. For differences in band energies we use the abbreviation

$$E_{ij} = E_i - E_j. \tag{10}$$

The energy value of interest, i.e. the energy at which we calculate the density of states, for example, is denominated by E. We now list the relevant formulas for calculating the weights and their derivatives.

In case the energy of interest is lower than or equal to the lowest energy in the tetrahedron ($E \leq E_1$), the tetrahedron does not contribute to the integral:

$$w_j = 0, \quad j = 1, 2, 3, 4 \tag{11}$$

In the case $E_1 < E < E_2$ we get a set of formulas:

$$C = \frac{V_T}{4V_G} \frac{(E-E_1)^3}{E_{21}E_{31}E_{41}} \tag{12a}$$

$$\frac{\partial C}{\partial E} = \frac{3V_T}{4V_G} \frac{(E-E_1)^2}{E_{21}E_{31}E_{41}} \tag{12b}$$

$$w_1 = C\Big[4 - (E-E_1)\Big(\frac{1}{E_{21}} + \frac{1}{E_{31}} + \frac{1}{E_{41}}\Big)\Big] \tag{12c}$$

$$\frac{\partial w_1}{\partial E} = \frac{\partial C}{\partial E}\Big[4 - (E-E_1)\Big(\frac{1}{E_{21}} + \frac{1}{E_{31}} + \frac{1}{E_{41}}\Big)\Big] - C\Big(\frac{1}{E_{21}} + \frac{1}{E_{31}} + \frac{1}{E_{41}}\Big) \tag{12d}$$

$$w_j = C\frac{E-E_1}{E_{j1}}, \quad j = 2, 3, 4 \tag{12e}$$

$$\frac{\partial w_j}{\partial E} = \frac{\partial C}{\partial E}\frac{E-E_1}{E_{j1}} + \frac{C}{E_{j1}} \tag{12f}$$

In the case $E_2 < E < E_3$ another set of formulas is used:

$$C_1 = \frac{V_T}{4V_G} \frac{(E-E_1)^2}{E_{31}E_{41}} \tag{13a}$$

$$\frac{\partial C_1}{\partial E} = \frac{V_T}{2V_G} \frac{E-E_1}{E_{31}E_{41}} \tag{13b}$$

$$C_2 = \frac{V_T}{4V_G} \frac{(E-E_1)(E-E_2)(E_3-E)}{E_{31}E_{32}E_{41}} \tag{13c}$$

$$\frac{\partial C_2}{\partial E} = \frac{V_T}{4V_G} \frac{1}{E_{31}E_{32}E_{41}}\Big[(E-E_1)(E_2-E) + (E-E_1)(E_3-E) + (E-E_2)(E_3-E)\Big] \tag{13d}$$

$$C_3 = \frac{V_T}{4V_G} \frac{(E-E_2)^2(E_4-E)}{E_{32}E_{41}E_{42}} \tag{13e}$$

$$\frac{\partial C_3}{\partial E} = \frac{V_T}{4V_G} \frac{E-E_2}{E_{32}E_{41}E_{42}}\Big[2(E_4-E) + (E_2-E)\Big] \tag{13f}$$

$$w_1 = C_1 + (C_1+C_2)\frac{E_3-E}{E_{31}} + (C_1+C_2+C_3)\frac{E_4-E}{E_{41}} \tag{13g}$$

$$\frac{\partial w_1}{\partial E} = \frac{\partial C_1}{\partial E} + \Big(\frac{\partial C_1}{\partial E} + \frac{\partial C_2}{\partial E}\Big)\frac{E_3-E}{E_{31}} - \frac{C_1+C_2}{E_{31}} + \Big(\frac{\partial C_1}{\partial E} + \frac{\partial C_2}{\partial E} + \frac{\partial C_3}{\partial E}\Big)\frac{E_4-E}{E_{41}} - \frac{C_1+C_2+C_3}{E_{41}} \tag{13h}$$

$$w_2 = C_1 + C_2 + C_3 + (C_2+C_3)\frac{E_3-E}{E_{32}} + C_3\frac{E_4-E}{E_{42}} \tag{13i}$$

$$\frac{\partial w_2}{\partial E} = \frac{\partial C_1}{\partial E} + \frac{\partial C_2}{\partial E} + \frac{\partial C_3}{\partial E} + \Big(\frac{\partial C_2}{\partial E} + \frac{\partial C_3}{\partial E}\Big)\frac{E_3-E}{E_{32}} - \frac{C_2+C_3}{E_{32}} + \frac{\partial C_3}{\partial E}\frac{E_4-E}{E_{42}} - \frac{C_3}{E_{42}} \tag{13j}$$

$$w_3 = (C_1+C_2)\frac{E-E_1}{E_{31}} + (C_2+C_3)\frac{E-E_2}{E_{32}} \tag{13k}$$

$$\frac{\partial w_3}{\partial E} = \Big(\frac{\partial C_1}{\partial E} + \frac{\partial C_2}{\partial E}\Big)\frac{E-E_1}{E_{31}} + \frac{C_1+C_2}{E_{31}} \tag{13l}$$

$$+\Big(\frac{\partial C_2}{\partial E}+\frac{\partial C_3}{\partial E}\Big)\frac{E-E_2}{E_{32}}+\frac{C_1+C_3}{E_{32}}$$

$$w_4=(C_1+C_2+C_3)\frac{E-E_1}{E_{41}}+C_3\frac{E-E_2}{E_{42}} \tag{13m}$$

$$\frac{\partial w_4}{\partial E}=\Big(\frac{\partial C_1}{\partial E}+\frac{\partial C_2}{\partial E}+\frac{\partial C_3}{\partial E}\Big)\frac{E-E_1}{E_{41}} \tag{13n}$$
$$+\frac{C_1+C_2+C_3}{E_{41}}+\frac{\partial C_3}{\partial E}\frac{E-E_2}{E_{42}}+\frac{C_3}{E_{42}}$$

The case $E_3 < E < E_4$ is given by:

$$C=\frac{V_T}{4V_G}\frac{(E_4-E)^3}{E_{41}E_{42}E_{43}} \tag{14a}$$

$$\frac{\partial C}{\partial E}=-\frac{3V_T}{4V_G}\frac{(E_4-E)^2}{E_{41}E_{42}E_{43}} \tag{14b}$$

$$w_j=\frac{V_T}{4V_G}-C\frac{E_4-E}{E_{4j}},\quad j=1,2,3 \tag{14c}$$

$$\frac{\partial w_j}{\partial E}=-\frac{\partial C}{\partial E}\frac{E_4-E}{E_{4j}}+\frac{C}{E_{4j}} \tag{14d}$$

$$w_4=\frac{V_T}{4V_G} \tag{14e}$$
$$-C\Big[4-\Big(\frac{1}{E_{41}}+\frac{1}{E_{42}}+\frac{1}{E_{43}}\Big)(E_4-e)\Big]$$

$$\frac{\partial w_4}{\partial E}=-\frac{\partial C}{\partial E}\Big[4-\Big(\frac{1}{E_{41}}+\frac{1}{E_{42}}+\frac{1}{E_{43}}\Big)(E_4-e)\Big] \tag{14f}$$
$$-C\Big(\frac{1}{E_{41}}+\frac{1}{E_{42}}+\frac{1}{E_{43}}\Big)$$

And finally, in case the energy of interest lies above the energy range of the tetrahedron ($E_4 < E$), the weights on the tetrahedron are constant:

$$w_j=\frac{V_T}{4V_G},\quad \frac{\partial w_j}{\partial E}=0,\quad j=1,2,3,4 \tag{15}$$

III. GPU TETRAHEDRON ALGORITHM AND IMPLEMENTATION CHOICES

A. Design goals and hardware constraints

GPU hardware differs from usual CPUs in two important ways: (i) A GPU contains many more compute cores than a CPU, usually thousands versus a few. (ii) Memory bandwidth within the GPU is much larger than between a CPU and normal memory.

Using the specialized hardware in the GPU comes at the expense of having to transfer data from the memory the CPU has direct access to, usually called *host* memory, to the memory the GPU has direct access to, usually called *device* memory. The goal of any GPU algorithm is, therefore, to use the higher memory and compute throughput of the GPU to generate large speedups that by far outweigh the overheads introduced by data transfer and synchronization between host and device code.

The usual parallelization strategy for CPU programs is to find a variable in which a formula is naturally parallel, and execute the associated program loop distributed across the available compute cores. While this strategy also works for simple algorithms when porting host code to the GPU, it is not applicable to the tetrahedron algorithm for Brillouin zone integration, as we will demonstrate later on.

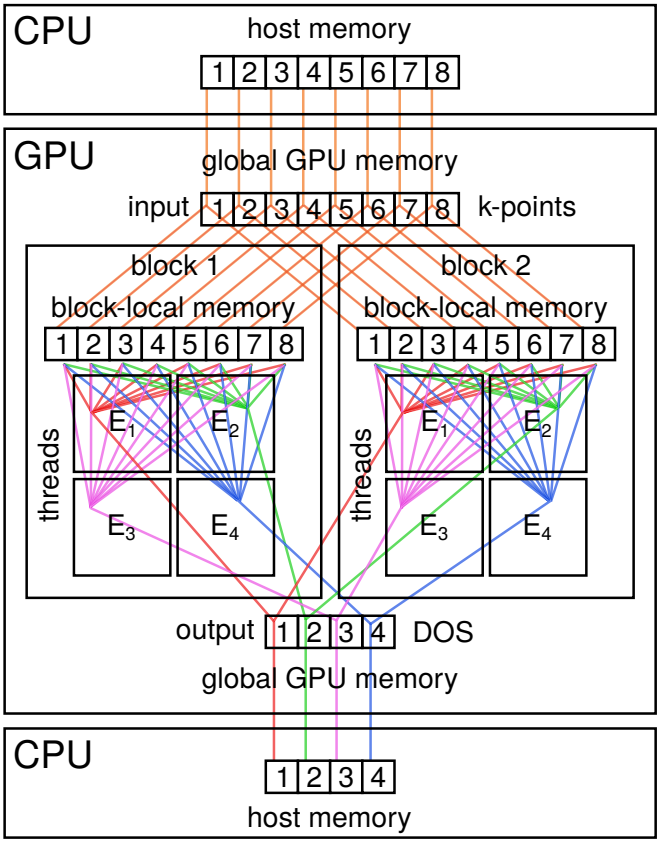

FIG. 3. (Color online) Visualization of memory access routes in our algorithm. Cooperation of compute threads organized into blocks on common chunks of data is key to achieving high performance.

The reason for this is that the integration is formulated as a sum of contributions of tetrahedra, which share corners. This leads to rather complex memory access patters, where points on the momentum grid are accessed several times. Feeding all available compute cores with enough data, therefore, becomes tricky, since the memory bandwidth even of the GPU is quickly exhausted when all compute cores work independent from each other like in the usual CPU parallelization strategy.

Fortunately, compute cores within the CUDA framework can cooperate on shared chunks of memory if they are organized into so-called *blocks*. Using these blocks is essential in memory-bound problems, since block-local memory can be filled from global GPU memory once and then be accessed cheaply by all compute cores in the block. Memory access within our algorithm is visualized in Fig. 3.

We first transfer all input data from host memory to GPU memory. Then we fetch the data that a block

TABLE I. Corner indices m_j belonging to the tetrahedron with index i within an elemental cube of the momentum grid. Compare Fig. 2.

i	m_1	m_2	m_3	m_4
1	1	2	3	6
2	2	3	4	6
3	1	3	5	6
4	3	4	6	8
5	3	5	6	7
6	3	6	7	8

of threads is working on to block-local memory. Subsequently the computational kernel is executed, in which threads solely read from block-local memory and accumulate results in thread-local memory (not explicitly shown in Fig. 3). These two are the key ingredients of the speedups we generate over a CPU implementation. After the computation is finished, results are written out to global GPU memory. An accumulation step within each block is omitted, since the results of each thread are independent in our implementation. After the results from each block have been returned, the final result accumulated over all blocks is copied to host memory.

B. Algorithm outline and implementation decisions

The algorithm for calculating the orbital-resolved density of states $\rho_m(E)$ (Eq. 9) can be broken down as follows: To every elemental cube on the momentum grid we assign a CUDA block with m threads. The number of external energies E for which the density of states is evaluated within this block is equal to m. That means, all threads within the block work on the same set of data, but at different energies of interest. We then loop over all six tetrahedra within the elemental cube. In a further inner loop we iterate over the band index. We calculate the integration weights for this band and finally calculate the orbital-resolved density of states in a loop over orbitals.

The definition of the tetrahedra in terms of corner indices of an elemental cube in the momentum grid is given in Table I. We put this index table into fast read-only memory, since it is a constant of our algorithm.

In the following we explain decisions we made to arrive at an efficient implementation. Since storing multidimensional arrays on GPUs is not possible without further complication, we assume band energies and orbital dependent matrix elements as a function of momentum are stored in linearized arrays.

Therefore, on the GPU we precalculate the indices of the eight corners of every elemental cube and store them. Subsequently, for every tetrahedron we sort its four energies in ascending order and store the corresponding indices of the original array, so that we can later find the matrix elements that belong to a certain band energy. Sorting the energies within a tetrahedron is necessary in order to know which formula for the integration weights has to be used. In principle these two steps can be executed within the main computational loop, however, this would mean that index calculation and energy sorting would be performed unnecessarily by every thread in the block, although indices and energy ordering are constant per cube.

Instead we wrote a separate computational kernel, to be executed before the main kernel, which does index precalculation and energy sorting without wasting any computational time. Splitting a computational problem into parts that require different parallel setups is often necessary to fully optimize a GPU algorithm.

The loop over energies of interest E is done in the following way: each energy is assigned a thread within the block. This way, it is guaranteed that threads within a block will never compete for write acess to memory. Avoiding memory collisions is essential in avoiding wait times that slow down the computation.

However, calculations may be carried out for several cubes in parallel, which could lead to write conflicts on the central output array in global GPU memory, where the final density of states is accumulated (see Fig. 3). Therefore, we first accumulate the results at thread-local level and then write the result out to the central result array using *atomic* operations, which mediate between threads competing for write access at the same memory locations at the expense of some management overhead.

Leaving out the thread-local accumulation of results before writing them out to global GPU memory leads to a severe decrease in performance, since the number of write collisions is roughly multiplied by the number of bands in the problem times six (the number of tetrahedra per cube), which results in significantly increased synchronization effort.

We have also implemented the possibility to let every CUDA block work on more than one elemental cube. In practice varying the number of cubes per block did not result in significantly altered performance.

Finally, we remark that all calculations are performed in single precision, since Nvidia consumer cards do not deliver maximum performance at double precision. When specialized GPUs for scientific computations are available, one can in principle switch to double precision. However, such high precision is actually rarely needed in solid state applications.

C. Benchmark method

We benchmark our GPU algorithm for the orbital-resolved density of states against an optimized equivalent CPU implementation. We use momentum grids with a different number of elemental cubes, i.e. a different resolution, to scale the workload. The smallest grid we investigate is a $10\times10\times10 = 1000$ k-point grid, the largest grid we investigate consists of $60 \times 60 \times 10 = 36000$ k-points. The grid is only refined in two dimensions, since the elec-

tronic structure of FeSe is almost two-dimensional. However, the algorithm run time is only sensitive to the total number of k-points, not their exact spatial distribution.

We evaluate the orbital-resolved density of states on an equidistantly spaced grid of 1024 energies. The quantities we calculate are the orbital-resolved density of states for each of the five Fe $3d$ orbitals (although two of them are actually degenerate) and the total density of states.

First, the eigenenergies and eigenvectors of the tight-binding Hamiltonian are calculated on the CPU. A standard tetrahedron code for the total density of states would only require eigenenergies. However, since we are interested also in the orbital-resolved density of states, we also store the matrix elements $a_n^m(\vec{k})$. Note that, although we here prepare the inputs for our code using a tight-binding formalism, we actually do not make any assumptions about the origin of input data. These could as well have been prepared using a density functional theory calculation.

As a reference point for measuring the speedup of the GPU algorithm, we measure the execution time of the tetrahedron integration on the CPU by taking the average execution time of five subsequent runs. The calculation time for preparing the inputs for the tetrahedron integration, i.e. diagonalizing the Hamiltonian, is excluded here, so that the CPU execution time we measure purely represents the effort invested in the tetrahedron integration.

For measuring the GPU execution time we investigate two different situations: first we measure the GPU execution time including the CUDA runtime setup and the data transfer from host memory to the GPU and back from the GPU to host memory. This situation reprents a worst case scenario, in which our GPU integration algorithm is integrated in an existing CPU code, where additional time has to be invested in copying data back and forth. Therefore, we refer to this situation as the *dirty* case. The execution time is measured in seven subsequent runs and we take the average of the five shortest run times to be the average run time. The reason behind this procedure is that consumer GPUs usually do not run at maximum clock speeds if they are not under load. Therefore, the first few runs are executed with the hardware clocked below its maximum capacity.

Second, we measure purely the execution time of the computational kernels on the GPU, without any CUDA runtime setup and without data transfer times to or from host memory. We refer to this situation as the *clean* case, which represents a best case scenario in which our algorithm is used in conjunction with an existing GPU code, where input data naturally reside in GPU memory already. Of course this implies that the dirty execution time is always larger than the clean execution time. The speedup factors are finally calculated from the average dirty/clean GPU execution time divided by the average CPU execution time.

The CPU execution time is naturally measured in a clean situation, since our inputs were initially prepared on the CPU and are readily available. Therefore, the *dirty speedup* is the lower bound for the speedup of our algorithm compared to a CPU implementation, while the *clean speedup* is the upper bound.

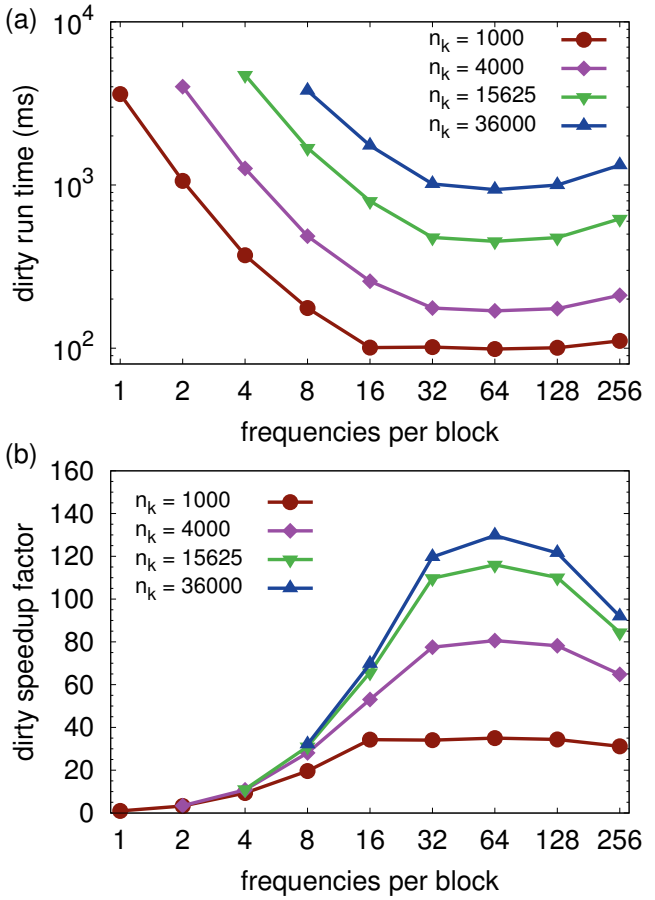

FIG. 4. (Color online) (a) Run time in milli-seconds for the GPU algorithm in the dirty scenario for different number of momentum grid points n_k. (b) Speedup factor in the dirty scenario relative to the CPU implementation.

The hardware we tested our algorithm on employs an Intel Core i7-4790K CPU and a Nvidia GeForce 970GTX GPU. The system runs on Ubuntu Linux 16.04 with CUDA runtime version 8.0 installed.

D. Benchmark results

We vary the number of frequencies per CUDA block to find the optimal speedup for the test case we defined in the previous section. The run times and speedup factors for the dirty case are shown in Fig. 4, while the results for the clean case are shown in Fig. 5.

The run time increase between different work load sizes is roughly linear, as expected. We, however, are mostly interested in the dependence of execution time on the number of frequencies per block, which controls the usage of block-local memory. Fig. 4(b) shows that the optimal

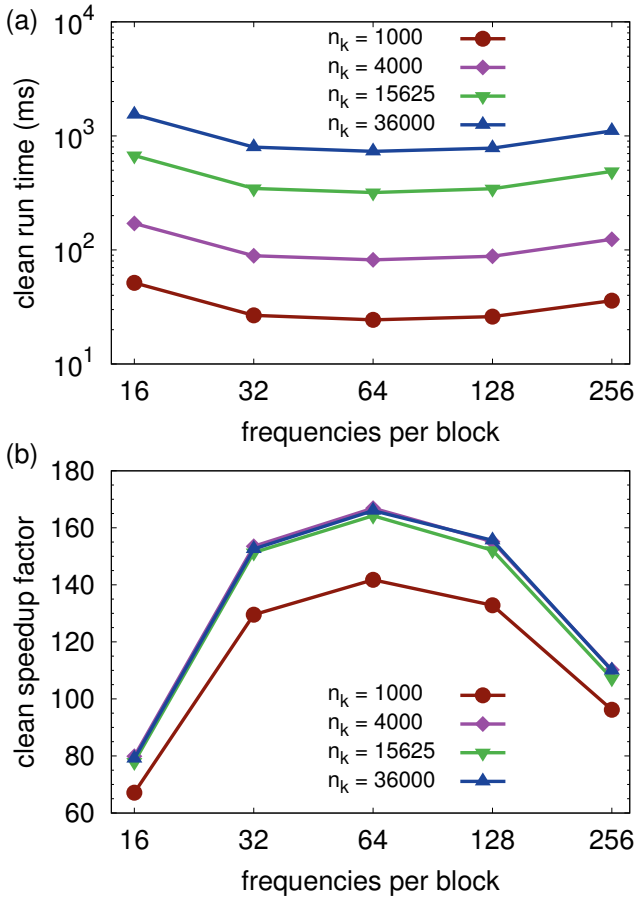

FIG. 5. (Color online) (a) Run time in milli-seconds for the GPU algorithm in the clean scenario for different number of momentum grid points n_k. (b) Speedup factor in the clean scenario relative to the CPU implementation.

number of frequencies per block is 64, independent of the work load size. Smaller and larger numbers of frequencies per CUDA block lead to an increase in execution time. In the dirty case the speedup also quite strongly depends on the size of the workload. Furthermore, for the smallest workload $n_k = 1000$ the run time is almost constant in between of 16 and 128 frequencies per block, which points to a significant overhead from data transfers.

Therefore, we turn to the clean execution times and speedups shown in Fig. 5. Here, all graphs show similar behavior, independent of the workload size. For sufficiently large workload sizes, the clean speedup is practically identical in all cases. Only for the smallest input data set we find a slightly reduced speedup, because the workload is too small to put the GPU under full load.

However, also the speedup graph for the smallest workload has a clear dome shape with the maximum at 64 frequencies per CUDA block. Since the clean run time in this case is much smaller than the dirty run time, the constant behavior in Fig. 4 simply means that the dirty run time mostly consists of data transfer time, while actual calculation time is negligible. Note, however, that

TABLE II. Run time and speedup factors as a function of the number of points on the momentum grid n_k. For the GPU runs we distinguish the settings dirty (GPUd) and clean (GPUc). For the GPU runs we show the results at the optimal setting of 64 frequencies per block. The speedup is calculated relative to the CPU run time.

n_k	CPU (ms)	GPUd (ms)	speedup	GPUc (ms)	speedup
1000	3455	98.7	35.0	24.4	141.8
4000	13652	169.4	80.6	81.8	166.9
15625	52370	451.4	116.0	318.9	164.2
36000	121515	936.7	129.7	732.0	166.0

even in this case our GPU algorithm performs about 35 times faster than a comparable CPU algorithm. Numerical values for run times and speedups are also given in Table II.

The optimal number of frequencies per block, which we determined, reflects that the speedup we achieved is a compromise between fast memory access across collaborating threads on the GPU, which we increasingly exploit as the number of frequencies per block grows, and the increasing effort for synchronizing these threads as block sizes grow. The tiny speedup factors we got in case of small number of frequencies per block illustrate that it is absolutely necessary to utilize thread collaboration on GPUs to achive significant advantages compared to CPU implementations. The usual CPU parallelization strategy, where threads work independent from each other, corresponds to the case of one frequency per block, i.e. where practically no speedup is achieved on GPU hardware.

It is possible that the optimal number of frequencies per block is different on different GPUs. However, as the speedup domes we observed are very wide, a trial-and-error search for the optimal settings on different hardware should not be difficult. Of course, GPUs with larger computational capacity may reach the optimal speedup as a function of workload size later than we did here.

IV. SUMMARY

We reviewed the tetrahedron method for Brillouin zone integrations and extended it to the case of the orbital-resolved density of states, supplying formulas for the momentum dependent weights within each tetrahedron.

Furthermore, we presented an efficient algorithm for calculating Brillouin zone integrals on modern graphics processing units and tested it on the problem of calculating the orbital-resolved density of states of an iron-based superconductor. Using the CUDA programming language we achieved an implementation that delivers large speedups of up to a factor ~ 165 compared to an analogous CPU implementation. We showed that optimization of memory access patterns was key to achieving these performance improvements.

Future work should concentrate on applying our algo-

rithm to computationally expensive quanitities that require Brillouin zone integrations, such as the susceptibilities employed in various many-body techniques.

[1] P. E. Blöchl, O. Jepsen, and O. K. Andersen, *Improved tetrahedron method for Brillouin-zone integrations*, Phys. Rev. B **49**, 16223 (1994).

[2] T. Fujiwara, S. Yamamoto, and Y. Ishii, *Generalization of the Iterative Perturbation Theory and Metal-Insulator Transition in Multi-Orbital Hubbard Bands*, J. Phys. Soc. Jpn. **72**, 777 (2003).

[3] J. Rath and A. J. Freeman, *Generalized magnetic susceptibilities in metals: Application of the analytic tetrahedron linear energy method to Sc*, Phys. Rev. **11**, 2109 (1975).

[4] Ch. Heil, H. Sormann, L. Boeri, M. Aichhorn, and W. von der Linden, *Accurate bare susceptibilities from full-potential ab initio calculations*, Phys. Rev. B **90**, 115143 (2014).

[5] S. Y. Savrasov, *Linear-Response Calculations of Lattice Dynamics Using Muffin-Tin Basis Sets*, Phys. Rev. Lett. **69**, 2819 (1992).

[6] S. Y. Savrasov and D. Y. Savrasov, *Electron-phonon interactions and related physical properties of metals from linear-response theory*, Phys. Rev. B **54**, 16487 (1996).

[7] M. Kawamura, Y. Gohda, and S. Tsuneyuki, *Improved tetrahedron method for the Brillouin-zone integration applicable to response functions*, Phys. Rev. B **89**, 094515 (2014).

[8] O. K. Andersen and L. Boeri, *On the multi-orbital band structure and itinerant magnetism of iron-based superconductors*, Ann. Phys. (Berlin) **523**, 8 (2011).

[9] H. Hosono and K. Kuroki, *Iron-based superconductors: Current status of materials and pairing mechanism*, Physica C **514**, 399 (2015).

[10] D. Guterding, H. O. Jeschke, P. J. Hirschfeld, and R. Valentí, *Unified picture of the doping dependence of superconducting transition temperatures in alkali metal/ammonia intercalated FeSe*, Phys. Rev. B **91**, 041112(R) (2015).