A fast Sparse SCF implementation: Application to DFTB

A. Scemama¹, M. Rapacioli¹

¹Laboratoire de Chimie et Physique Quantiques / IRSAMC, Toulouse, France

25 June 2013





- 2 Preparation: partition of the 3D space
- Initial Guess of Molecular Orbitals (MOs)
- Orthonormalization of MOs
- 5 Partial Diagonalization of ${\cal H}$

6 Results

- Floating-point operations (flops) scale as $\mathcal{O}(N^3)$
- Memory accesses scale as $\mathcal{O}(N^2)$
- Memory accesses have a very high latency (>300 CPU cycles) (1 RAM access is 50–90 nanoseconds: 130–230 CPU cycles at 2.6GHz) but
- Regular memory access patterns can be pre-fetched by the CPU in the caches, hiding the memory latencies
- Compute-bound, very efficient

- Floating-point operations (flops) scale as $\mathcal{O}(N^3)$
- Memory accesses scale as $\mathcal{O}(N^2)$
- Memory accesses have a very high latency (>300 CPU cycles) (1 RAM access is 50–90 nanoseconds: 130–230 CPU cycles at 2.6GHz) but
- Regular memory access patterns can be pre-fetched by the CPU in the caches, hiding the memory latencies
- Compute-bound, very efficient

- Floating-point operations (flops) scale as $\mathcal{O}(N^3)$
- Memory accesses scale as $\mathcal{O}(N^2)$
- Memory accesses have a very high latency (>300 CPU cycles) (1 RAM access is 50–90 nanoseconds: 130–230 CPU cycles at 2.6GHz) but
- Regular memory access patterns can be pre-fetched by the CPU in the caches, hiding the memory latencies
- Compute-bound, very efficient

- Floating-point operations (flops) scale as $\mathcal{O}(N^3)$
- Memory accesses scale as $\mathcal{O}(N^2)$
- Memory accesses have a very high latency (>300 CPU cycles) (1 RAM access is 50–90 nanoseconds: 130–230 CPU cycles at 2.6GHz) but
- Regular memory access patterns can be pre-fetched by the CPU in the caches, hiding the memory latencies
- Compute-bound, very efficient

- Floating-point operations (flops) scale as $\mathcal{O}(N^3)$
- Memory accesses scale as $\mathcal{O}(N^2)$
- Memory accesses have a very high latency (>300 CPU cycles) (1 RAM access is 50–90 nanoseconds: 130–230 CPU cycles at 2.6GHz) but
- Regular memory access patterns can be pre-fetched by the CPU in the caches, hiding the memory latencies
- Compute-bound, very efficient

Sparse \times sparse matrix products are very difficult to perform efficiently:

- Memory accesses scale as $\mathcal{O}(N)$
- Floating-point operations (flops) scales as $\mathcal{O}(N)$
- Necessarily memory-bound

- Compiler can't produce efficient code (memory indirections, branches)
- No vectorization

Sparse \times sparse matrix products are very difficult to perform efficiently:

- Memory accesses scale as $\mathcal{O}(N)$
- Floating-point operations (flops) scales as $\mathcal{O}(N)$
- Necessarily memory-bound

- Compiler can't produce efficient code (memory indirections, branches)
- No vectorization

Sparse \times sparse matrix products are very difficult to perform efficiently:

- Memory accesses scale as $\mathcal{O}(N)$
- Floating-point operations (flops) scales as $\mathcal{O}(N)$
- Necessarily memory-bound

- Compiler can't produce efficient code (memory indirections, branches)
- No vectorization

Sparse \times sparse matrix products are very difficult to perform efficiently:

- Memory accesses scale as $\mathcal{O}(N)$
- Floating-point operations (flops) scales as $\mathcal{O}(N)$
- Necessarily memory-bound

- Compiler can't produce efficient code (memory indirections, branches)
- No vectorization

Sparse \times sparse matrix products are very difficult to perform efficiently:

- Memory accesses scale as $\mathcal{O}(N)$
- Floating-point operations (flops) scales as $\mathcal{O}(N)$
- Necessarily memory-bound

- Compiler can't produce efficient code (memory indirections, branches)
- No vectorization

- All arrays are 256-bit aligned (compiler directives)
- Leading dimensions are multiples of 8 : all columns 256-bit aligned
- Unroll and jam to reduce nb of stores
- All inner-most loops are fully vectorized : always a multiple of 8 loop cycles
- Static analysis of the binary (MAQAO): All inner-most loops can reach the peak of 16 flops/cycle
- "Only" 60%: Both memory and flops $\mathcal{O}(N^2)$
- Memory bound, but memory latencies reduced by pre-fetching in the dense matrix.

- All arrays are 256-bit aligned (compiler directives)
- Leading dimensions are multiples of 8 : all columns 256-bit aligned
- Unroll and jam to reduce nb of stores
- All inner-most loops are fully vectorized : always a multiple of 8 loop cycles
- Static analysis of the binary (MAQAO): All inner-most loops can reach the peak of 16 flops/cycle
- "Only" 60%: Both memory and flops $\mathcal{O}(N^2)$
- Memory bound, but memory latencies reduced by pre-fetching in the dense matrix.

- All arrays are 256-bit aligned (compiler directives)
- Leading dimensions are multiples of 8 : all columns 256-bit aligned
- Unroll and jam to reduce nb of stores
- All inner-most loops are fully vectorized : always a multiple of 8 loop cycles
- Static analysis of the binary (MAQAO): All inner-most loops can reach the peak of 16 flops/cycle
- "Only" 60%: Both memory and flops $\mathcal{O}(N^2)$
- Memory bound, but memory latencies reduced by pre-fetching in the dense matrix.

- All arrays are 256-bit aligned (compiler directives)
- Leading dimensions are multiples of 8 : all columns 256-bit aligned
- Unroll and jam to reduce nb of stores
- All inner-most loops are fully vectorized : always a multiple of 8 loop cycles
- Static analysis of the binary (MAQAO): All inner-most loops can reach the peak of 16 flops/cycle
- "Only" 60%: Both memory and flops $\mathcal{O}(N^2)$
- Memory bound, but memory latencies reduced by pre-fetching in the dense matrix.

- All arrays are 256-bit aligned (compiler directives)
- Leading dimensions are multiples of 8 : all columns 256-bit aligned
- Unroll and jam to reduce nb of stores
- All inner-most loops are fully vectorized : always a multiple of 8 loop cycles
- Static analysis of the binary (MAQAO): All inner-most loops can reach the peak of 16 flops/cycle
- "Only" 60%: Both memory and flops $\mathcal{O}(N^2)$
- Memory bound, but memory latencies reduced by pre-fetching in the dense matrix.

- All arrays are 256-bit aligned (compiler directives)
- Leading dimensions are multiples of 8 : all columns 256-bit aligned
- Unroll and jam to reduce nb of stores
- All inner-most loops are fully vectorized : always a multiple of 8 loop cycles
- Static analysis of the binary (MAQAO): All inner-most loops can reach the peak of 16 flops/cycle
- "Only" 60%: Both memory and flops $\mathcal{O}(N^2)$
- Memory bound, but memory latencies reduced by pre-fetching in the dense matrix.

- All arrays are 256-bit aligned (compiler directives)
- Leading dimensions are multiples of 8 : all columns 256-bit aligned
- Unroll and jam to reduce nb of stores
- All inner-most loops are fully vectorized : always a multiple of 8 loop cycles
- Static analysis of the binary (MAQAO): All inner-most loops can reach the peak of 16 flops/cycle
- "Only" 60%: Both memory and flops $\mathcal{O}(N^2)$
- Memory bound, but memory latencies reduced by pre-fetching in the dense matrix.

Sparse representation of the matrices

Each sparse matrix is represented by two arrays:

- An array of non-zero indices per column
- An array of corresponding values

Example: One column of matrix V





- Represent one sparse matrix as a collection of small dense sub-matrices
- Each dense sub-matrix is 256-bit aligned
- Each column of the sub-matrix is 256-bit aligned (padding)
- The number of loop-cycles is set to a multiple of 8

- Represent one sparse matrix as a collection of small dense sub-matrices
- Each dense sub-matrix is 256-bit aligned
- Each column of the sub-matrix is 256-bit aligned (padding)
- The number of loop-cycles is set to a multiple of 8

- Represent one sparse matrix as a collection of small dense sub-matrices
- Each dense sub-matrix is 256-bit aligned
- Each column of the sub-matrix is 256-bit aligned (padding)
- The number of loop-cycles is set to a multiple of 8

- Represent one sparse matrix as a collection of small dense sub-matrices
- Each dense sub-matrix is 256-bit aligned
- Each column of the sub-matrix is 256-bit aligned (padding)
- The number of loop-cycles is set to a multiple of 8





- Preparation: partition of the 3D space
 - Initial Guess of Molecular Orbitals (MOs)
- Orthonormalization of MOs
- 5 Partial Diagonalization of ${\cal H}$

6 Results

- A set of *m* centers (*means*) is first distributed evenly in the 3D-space
- Each molecule is attached to its closest center, such that each center is connected to 4 molecules
- The position of the centers is moved to the centroid of the connected molecules
- Go back to step 2 until the partition doesn't change

- A set of *m* centers (*means*) is first distributed evenly in the 3D-space
- Each molecule is attached to its closest center, such that each center is connected to 4 molecules
- The position of the centers is moved to the centroid of the connected molecules
- Go back to step 2 until the partition doesn't change

- A set of *m* centers (*means*) is first distributed evenly in the 3D-space
- Each molecule is attached to its closest center, such that each center is connected to 4 molecules
- The position of the centers is moved to the centroid of the connected molecules
- Go back to step 2 until the partition doesn't change

- A set of *m* centers (*means*) is first distributed evenly in the 3D-space
- Each molecule is attached to its closest center, such that each center is connected to 4 molecules
- The position of the centers is moved to the centroid of the connected molecules
- Go back to step 2 until the partition doesn't change











Efficient sparse matrix products
Preparation: partition of the 3D space
Initial Guess of Molecular Orbitals (MOs)
Orthonormalization of MOs
Partial Diagonalization of ${\cal H}$
Results

- The atoms are then ordered by k-means centers
- *k-means* neighbours are centers with at least 2 atoms less than 20 a.u




- 2 Preparation: partition of the 3D space
- Initial Guess of Molecular Orbitals (MOs)
- Orthonormalization of MOs
- 5 Partial Diagonalization of ${\cal H}$

6 Results

- For each molecule, perform an independent non-SCC DFTB calculation (Solve $\mathcal{H}C = ESC$)
- Pack together occupied MOs of each k-means center
- Pack together virtual MOs of each k-means center

- MOs belonging to the same molecule are orthonormal
- MOs belonging to *k-means* centers which are not neighbours have a zero overlap
- MOs have non-zero coefficients only on basis functions which belong to the same molecule
- The orbitals are not orthonormal, but both the overlap (S) and the MO coefficient (C) matrix are sparse.

- For each molecule, perform an independent non-SCC DFTB calculation (Solve $\mathcal{H}C = ESC$)
- Pack together occupied MOs of each k-means center
- Pack together virtual MOs of each k-means center

- MOs belonging to the same molecule are orthonormal
- MOs belonging to k-means centers which are not neighbours have a zero overlap
- MOs have non-zero coefficients only on basis functions which belong to the same molecule
- The orbitals are not orthonormal, but both the overlap (S) and the MO coefficient (C) matrix are sparse.

- For each molecule, perform an independent non-SCC DFTB calculation (Solve $\mathcal{H}C = ESC$)
- Pack together occupied MOs of each k-means center
- Pack together virtual MOs of each *k-means* center

- MOs belonging to the same molecule are orthonormal
- MOs belonging to k-means centers which are not neighbours have a zero overlap
- MOs have non-zero coefficients only on basis functions which belong to the same molecule
- The orbitals are not orthonormal, but both the overlap (S) and the MO coefficient (C) matrix are sparse.

- For each molecule, perform an independent non-SCC DFTB calculation (Solve $\mathcal{H}C = ESC$)
- Pack together occupied MOs of each k-means center
- Pack together virtual MOs of each k-means center

- MOs belonging to the same molecule are orthonormal
- MOs belonging to *k-means* centers which are not neighbours have a zero overlap
- MOs have non-zero coefficients only on basis functions which belong to the same molecule
- The orbitals are not orthonormal, but both the overlap (S) and the MO coefficient (C) matrix are sparse.

- For each molecule, perform an independent non-SCC DFTB calculation (Solve $\mathcal{H}C = ESC$)
- Pack together occupied MOs of each k-means center
- Pack together virtual MOs of each k-means center

- MOs belonging to the same molecule are orthonormal
- MOs belonging to k-means centers which are not neighbours have a zero overlap
- MOs have non-zero coefficients only on basis functions which belong to the same molecule
- The orbitals are not orthonormal, but both the overlap (S) and the MO coefficient (C) matrix are sparse.

- For each molecule, perform an independent non-SCC DFTB calculation (Solve $\mathcal{H}C = ESC$)
- Pack together occupied MOs of each k-means center
- Pack together virtual MOs of each k-means center

- MOs belonging to the same molecule are orthonormal
- MOs belonging to k-means centers which are not neighbours have a zero overlap
- MOs have non-zero coefficients only on basis functions which belong to the same molecule
- The orbitals are not orthonormal, but both the overlap (S) and the MO coefficient (C) matrix are sparse.

- For each molecule, perform an independent non-SCC DFTB calculation (Solve $\mathcal{H}C = ESC$)
- Pack together occupied MOs of each k-means center
- Pack together virtual MOs of each k-means center

- MOs belonging to the same molecule are orthonormal
- MOs belonging to k-means centers which are not neighbours have a zero overlap
- MOs have non-zero coefficients only on basis functions which belong to the same molecule
- The orbitals are not orthonormal, but both the overlap (S) and the MO coefficient (C) matrix are sparse.

Outline

- Efficient sparse matrix products
- Preparation: partition of the 3D space
- Initial Guess of Molecular Orbitals (MOs)
- Orthonormalization of MOs
- 5 Partial Diagonalization of $\mathcal H$

6 Results

Diagonalization of $C^{\dagger}SC$

The C matrix is already stored sparse

- Compute S (sparse)
- Compute C[†]SC (sparse)
- In Normalize using diagonal elements
- Perform Jacobi-like rotations to remove the largest off-diagonal elements of C[†]SC
- Go back to step 3 until the largest off-diagonal element is below a threshold

Diagonalization of $C^{\dagger}SC$

- The C matrix is already stored sparse
- Ompute S (sparse)
- Compute C[†]SC (sparse)
- Ormalize using diagonal elements
- Perform Jacobi-like rotations to remove the largest off-diagonal elements of C[†]SC
- Go back to step 3 until the largest off-diagonal element is below a threshold

Diagonalization of $C^{\dagger}SC$

- The C matrix is already stored sparse
- Ompute S (sparse)
- **③** Compute $C^{\dagger}SC$ (sparse)
- Normalize using diagonal elements
- Perform Jacobi-like rotations to remove the largest off-diagonal elements of C[†]SC
- Go back to step 3 until the largest off-diagonal element is below a threshold

Diagonalization of $C^{\dagger}SC$

- The C matrix is already stored sparse
- Ompute S (sparse)
- **③** Compute $C^{\dagger}SC$ (sparse)
- Ormalize using diagonal elements
- Perform Jacobi-like rotations to remove the largest off-diagonal elements of C[†]SC
- Go back to step 3 until the largest off-diagonal element is below a threshold

Diagonalization of $C^{\dagger}SC$

- The C matrix is already stored sparse
- Ompute S (sparse)
- **③** Compute $C^{\dagger}SC$ (sparse)
- Ormalize using diagonal elements
- Perform Jacobi-like rotations to remove the largest off-diagonal elements of C[†]SC
- Go back to step 3 until the largest off-diagonal element is below a threshold

Diagonalization of $C^{\dagger}SC$

- The C matrix is already stored sparse
- Ompute S (sparse)
- **③** Compute $C^{\dagger}SC$ (sparse)
- Ormalize using diagonal elements
- Perform Jacobi-like rotations to remove the largest off-diagonal elements of C[†]SC
- Go back to step 3 until the largest off-diagonal element is below a threshold

Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOs) Orthoormalization of MOs

Partial Diagonalization of H Results

Parallel implementation of $C^{\dagger}SC$



Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOs) Orthonormalization of MOs

Partial Diagonalization of H Results

Parallel implementation of $C^{\dagger}SC$



Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOs) Orthonormalization of MOs

Partial Diagonalization of H Results

Parallel implementation of $C^{\dagger}SC$



Results

Parallel implementation of $C^{\dagger}SC$

Transpose C+SC



Parallel implementation of Jacobi-like rotations

- Pick the next rotation (i, j)
- If rotation (i, j) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate *i* and *j*
- Mark (*i*, *j*) as done
- Free locks i and j
- O Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

- Pick the next rotation (*i*, *j*)
- If rotation (*i*, *j*) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate *i* and *j*
- Mark (*i*, *j*) as done
- Free locks *i* and *j*
- O Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

- Pick the next rotation (*i*, *j*)
- If rotation (*i*, *j*) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate *i* and *j*
- Mark (*i*, *j*) as done
- Free locks *i* and *j*
- O Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

- Pick the next rotation (*i*, *j*)
- If rotation (i, j) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate i and j
- Mark (*i*, *j*) as done
- Free locks *i* and *j*
- O Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

- Pick the next rotation (*i*, *j*)
- If rotation (i, j) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate i and j
- Mark (*i*, *j*) as done
- Free locks *i* and *j*
- O Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

- **O** Pick the next rotation (i, j)
- If rotation (i, j) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate i and j
- Mark (*i*, *j*) as done
- Free locks *i* and *j*
- O Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

- Pick the next rotation (*i*, *j*)
- If rotation (i, j) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate i and j
- Mark (*i*, *j*) as done
- Free locks *i* and *j*
- O Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

- **O** Pick the next rotation (i, j)
- If rotation (i, j) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate i and j
- Mark (*i*, *j*) as done
- Free locks *i* and *j*
- O Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

First, dress the list of rotations to do: $|(\mathbf{C}^{\dagger}\mathbf{S}\mathbf{C})_{i>j}| > \epsilon$. Prepare a 1D-array of OpenMP locks, (one lock for each MO). All CPUs do at the same time:

- **O** Pick the next rotation (i, j)
- 2 If rotation (i, j) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate i and j
- Mark (*i*, *j*) as done
- Free locks *i* and *j*

Go back to step 1 until all rotations are done

Parallel implementation of Jacobi-like rotations

- **O** Pick the next rotation (i, j)
- 2 If rotation (i, j) is already done, go to step 1
- Try to take lock i
- If not possible, go to step 1
- Try to take lock j
- If not possible, free lock i and go to step 1
- Rotate i and j
- Mark (*i*, *j*) as done
- Free locks *i* and *j*
- Observe to the step 1 until all rotations are done

Outline

- Efficient sparse matrix products
- Preparation: partition of the 3D space
- Initial Guess of Molecular Orbitals (MOs)
- Orthonormalization of MOs
- 5 Partial Diagonalization of $\mathcal H$

6 Results

- The energy is independent of unitary transformations among occupied-occupied blocks and virtual-virtual blocks
- Brillouin's theorem: It is sufficient to zero the occupied-virtual blocks
- Avoiding occupied-occupied and virtual-virtual rotations keeps the orbitals local (sparse)
- The partial diagonalization of C[†]HC is performed like the diagonalization C[†]SC, computing only the occupied-virtual blocks.
- Double precision is not required here \longrightarrow Single precision.
- The self-consistence step is inserted before each C[†]HC product

- The energy is independent of unitary transformations among occupied-occupied blocks and virtual-virtual blocks
- Brillouin's theorem: It is sufficient to zero the occupied-virtual blocks
- Avoiding occupied-occupied and virtual-virtual rotations keeps the orbitals local (sparse)
- The partial diagonalization of C[†]HC is performed like the diagonalization C[†]SC, computing only the occupied-virtual blocks.
- Double precision is not required here \longrightarrow Single precision.
- The self-consistence step is inserted before each C[†]HC product

- The energy is independent of unitary transformations among occupied-occupied blocks and virtual-virtual blocks
- Brillouin's theorem: It is sufficient to zero the occupied-virtual blocks
- Avoiding occupied-occupied and virtual-virtual rotations keeps the orbitals local (sparse)
- The partial diagonalization of C[†]HC is performed like the diagonalization C[†]SC, computing only the occupied-virtual blocks.
- Double precision is not required here \longrightarrow Single precision.
- The self-consistence step is inserted before each C[†]HC product

- The energy is independent of unitary transformations among occupied-occupied blocks and virtual-virtual blocks
- Brillouin's theorem: It is sufficient to zero the occupied-virtual blocks
- Avoiding occupied-occupied and virtual-virtual rotations keeps the orbitals local (sparse)
- The partial diagonalization of C[†]HC is performed like the diagonalization C[†]SC, computing only the occupied-virtual blocks.
- Double precision is not required here \longrightarrow Single precision.
- The self-consistence step is inserted before each C[†]HC product

- The energy is independent of unitary transformations among occupied-occupied blocks and virtual-virtual blocks
- Brillouin's theorem: It is sufficient to zero the occupied-virtual blocks
- Avoiding occupied-occupied and virtual-virtual rotations keeps the orbitals local (sparse)
- The partial diagonalization of C[†]HC is performed like the diagonalization C[†]SC, computing only the occupied-virtual blocks.
- Double precision is not required here \longrightarrow Single precision.
- The self-consistence step is inserted before each C[†]HC product

- The energy is independent of unitary transformations among occupied-occupied blocks and virtual-virtual blocks
- Brillouin's theorem: It is sufficient to zero the occupied-virtual blocks
- Avoiding occupied-occupied and virtual-virtual rotations keeps the orbitals local (sparse)
- The partial diagonalization of C[†]HC is performed like the diagonalization C[†]SC, computing only the occupied-virtual blocks.
- Double precision is not required here \longrightarrow Single precision.
- The self-consistence step is inserted before each C[†]HC product
Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOS) Orthonormalization of MOS Partial Diagonalization of H Results

Outline

- Efficient sparse matrix products
- Preparation: partition of the 3D space
- Initial Guess of Molecular Orbitals (MOs)
- Orthonormalization of MOs
- 5 Partial Diagonalization of ${\cal H}$

6 Results

Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOs) Orthonormalization of MOs Partial Diagonalization of H Results

Benchmark: Boxes of water molecules, up to 120 000 atoms



- One compute node (OpenMP)
- Intel Xeon E5-2670, 2.7GHz (up to 3.3 GHz turbo)
- 2 sockets, 8 cores/socket
- 64Gb RAM

Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOs) Orthonormalization of MOs Partial Diagonalization of H Besults

Numerical precision

The relative error $\frac{E-E_{\text{LAPACK}}}{E_{\text{LAPACK}}}$ is below 10^{-9} :

#	LAPACK	This work	Relative Error
184	-749.639387 <mark>02</mark>	-749.639387 <mark>11</mark>	1.2 10 ⁻¹⁰
368	-1499.331634 <mark>92</mark>	-1499.331634 <mark>83</mark>	0.6 10 ⁻¹⁰
736	-2998.74237 <mark>166</mark>	-2998.74237 <mark>211</mark>	1.5 10 ⁻¹⁰
1472	-5997.7826 <mark>7988</mark>	-5997.7826 <mark>8084</mark>	1.6 10 ⁻¹⁰

The sparse cut-off criterion is adjusted with the target SCC convergence criterion ϵ :

our numerical errors are always lower than $\epsilon/10$

Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOs) Orthonormalization of MOs Partial Diagonalization of H Results

Total wall time



A. Scemama, M. Rapacioli Sparse DFTB

Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOS) Orthonormalization of MOS Partial Diagonalization of H Results

Parallel speed-up



A. Scemama, M. Rapacioli Sparse DFTB

Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOS) Orthonormalization of MOS Partial Diagonalization of H Results

Parallel efficiency



A. Scemama, M. Rapacioli Sparse DFTB

Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOs) Orthonormalization of MOs Partial Diagonalization of A

Results

Computational efficiency

- Orthonormalization: 18.5 GFlops/s DP (peak: 345.6 GFlops/s), L3: 83 GiB/sec (RAM peak: 102 GiB/sec)
- Partial diagonalization: 16.5 GFlops/s SP (peak: 691.2 GFlops/s), L3: 64 GiB/sec
- Memory latency -bound (50-90 ns)
- Memory consumption: ~ 1 MiB /water molecule

Efficient sparse matrix products Preparation: partition of the 3D space Initial Guess of Molecular Orbitals (MOS) Orthonormalization of MOS Partial Diagonalization of H Besults

More results

CALMIP SGI Altix UV: (48 sockets E7-8837 @ 2.67GHz, 8 cores/socket, 3Tb RAM) (NUMA, RAM latencies: 90– > 1000 ns) Results with 128 cores

- 184 000 molecules : 3.0 hours
- 244 904 molecules : 5.5 hours
- 317 952 molecules (1 million atoms!) : 9.0 hours