

I. SCALING INFORMATION

A. Performance/Scaling of the B-Spline Atomic R-Matrix Code

Figure 1 provides some scaling information for major parts of the BSR package. While BSR_BREIT and BSR_MAT scale well until about 800 cores in the example shown, the performance of BSR_HD deteriorates rapidly when the number of cores exceeds about 400-500. This lack of scaling efficiency in the BSR_HD step is due to; 1) the SCALAPACK [1] diagonalization, and 2) the ALL-TO-ALL operation of required to collect the eigenvectors from the distributed matrix. By simply collecting only the surface amplitudes, the scaling would be improved significantly. This will be implemented in the codes during the upcoming year.

Specifically, runs were performed on the Knight’s Landing (KNL) and Skylake (SKX) nodes on Stampede 2 with 256, 512, 768, and 1024 cores. We used 64 cores per node on the KNL nodes and 48 cores per node on the SKX nodes. On SuperMIC, 20 cores per node were used. If nodes rather than cores are a measure of performance then the KNL system on Stampede-2 and SuperMIC are comparable. The figure also shows that the SKX cores outperform the KNL cores by about a factor of three. This is disappointing but in line with what others have seen on the two systems. Unfortunately, we found that using more than one thread on the KNL cores slows the program down significantly. We currently do not have a clear explanation for this, but we intend to look into it in the next allocations period. Finally, while we realize that SuperMIC is no longer available to XSEDE, we trust that this information is still useful. It basically shows that we can expect to move rather easily to other machines, e.g., Comet and Bridges.

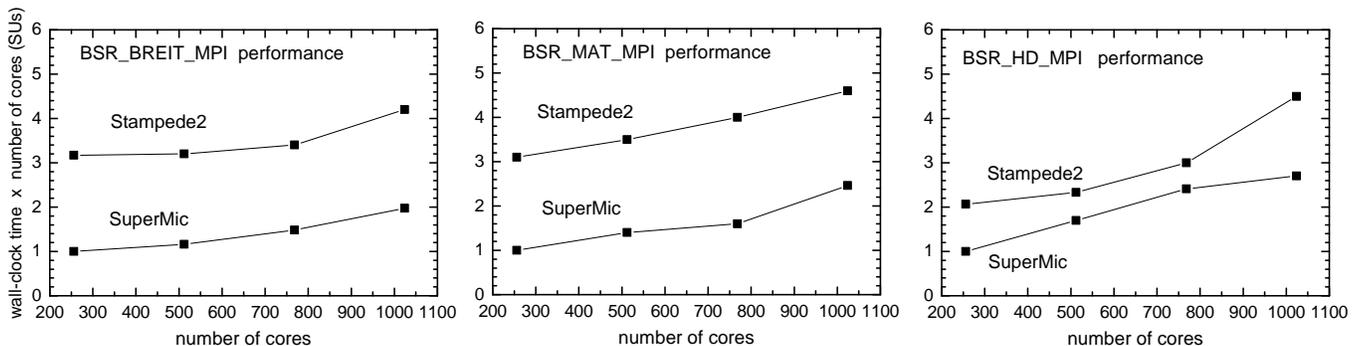


FIG. 1. The product of wall-time and the number of cores (normalized to the smallest number as a function of the latter) for BSR_BREIT, BSR_MAT, and BSR_HD on Stampede 2 and SuperMIC. The rank of the matrix was approximately 50,000.

We emphasize that in practice we do *not* waste SUs to the extent illustrated above, where doubling the number of cores from 512 to 1,024 increased the product of wall-time and the number of cores by a factor of $\approx 2 - 2.5$ on Stampede 2, i.e., it resulted in an increased wall-time as well. In fact, we always try to optimize the number of cores by using the following criteria:

- The matrix has to fit into the available memory, which sets the minimum number of cores.
- Increase the number of cores only to the extent that scaling does not deteriorate significantly while also ensuring that the maximum allowed wall-time is not exceeded.

This strategy ultimately sets the maximum size of the matrices that we can handle, and hence the size of the physical problems that we can tackle. To date, we have been able to deal with matrices of rank up to 400,000. Running BSR_HD represents the most expensive part of our work, and it has to be done for every partial-wave symmetry.

B. Performance/Scaling of the CCC Code

The CCC code has been parallelized using OpenMP on a node plus MPI between nodes. Recently, the code has been ported to GPU’s using OpenACC. The manner in which the calculation scales is shown in Fig. 2 for the two major computational steps in the code, formation of the CCC matrix and the solution of the large set of linear equations after the matrix has been formed. The first step, matrix formation, an essentially parallelizable operation, shows strong scaling, while the solution of the set of linear equations via ScaLAPACK degrades as the number of nodes increases. In addition, the GPU version of the code shows a performance enhancement of about 10 over the CPU version.

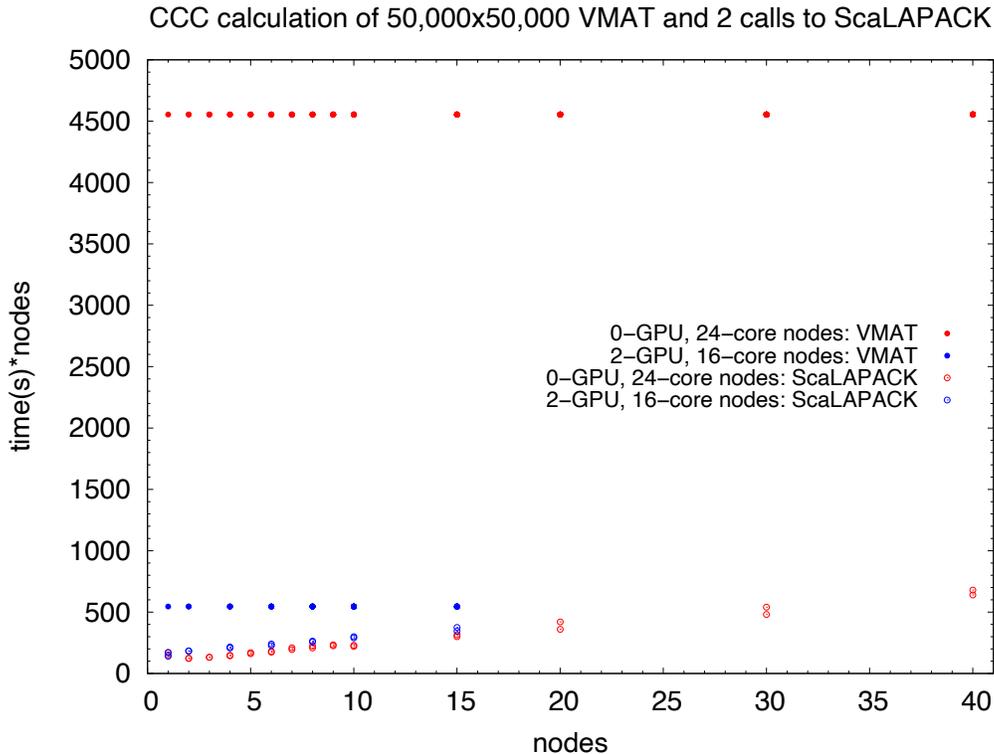


FIG. 2. The product of walltime and the number of cores (nodes) as a function of the latter for a CPU/GPU-based calculation on a 50,000 by 50,000 CCC matrix.

C. Performance/Scaling of the tRecX Code

Scaling is strongly application-dependent, mostly driven by the locality properties of the system’s representation. Performance in tRecX is achieved by a highly structured representation the operators. Most of the time is spent in time-propagation with the repeated application of block-sparse operators, where each block in itself is structured, e.g., as tensor products or singular-value representations. The linear dimension of the blocks varies from 10 to 400. The blocks determine the code’s granularity and are distributed based on self-measured CPU load. The distribution algorithms also attempts a trade-off between load-balance and communication. On the block level LAPACK and EIGEN routines are used. Expansion of the full operator into a matrix for the use of parallel linear-algebra software, while possibly improving formal scaling, ensues the loss of tensor structure and dramatic increase of float count.

Scaling data are provided in Fig. 3 for the calculations of Ref. [2], where most of the load is used in the electron interaction. The scaling data are for actual production-type runs including processing and output of intermediate results after each timestep, which involves non-local operations. It does not include setup times, as setup data can be reused among the runs of a parameter study. Results are for *strong scaling*, i.e., a problem of fixed size is distributed over an increasing number of codes with a corresponding increase of the relative role of communication.

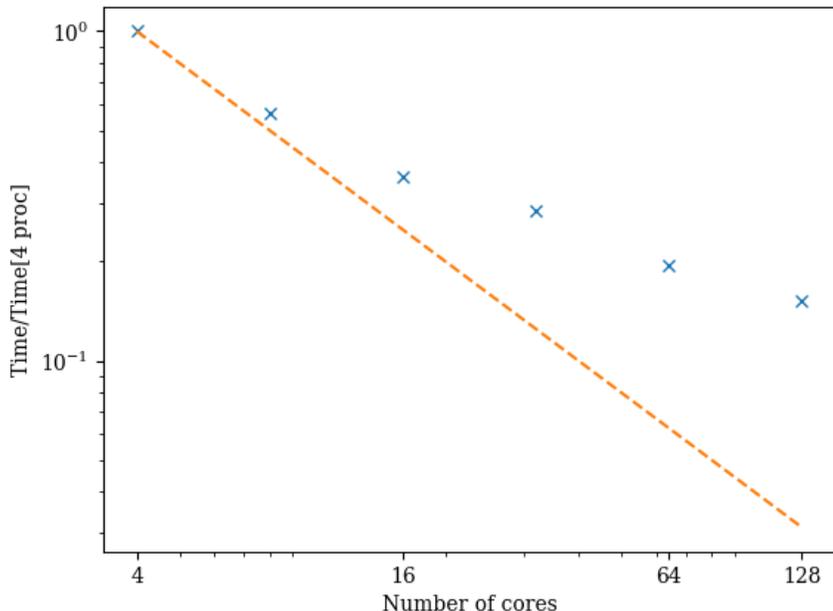


FIG. 3. Strong scaling of tRecX for double-ionization of Helium at laser wave length 400 nm. The results were obtained on the Munich Cool-MUC cluster. The specific section of Cool-MUC has nodes with two sockets with 16 cores each of an Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz with 192GB RAM DDR4@2666MHz per node. The dashed line is ideal scaling.

D. Performance/Scaling of the Xchem Code

The Xchem code is using the well-known OpenMolcas package [?] to provide the Gaussian integrals and its performance is highly dependent of the disk type. In general, using a SSD disk, the OpenMP parallelization is not very good as it can be seen in Fig ???. However, the code is using the SLURM job manager to highly parallelized between nodes. In this case, since there are several parts of the code where the jobs can be sent independently, we use the SLURM manager to send several jobs.

In general, the so-called local part (i.e. a regular quantum chemical calculation using OpenMolcas [?]) is evaluated separately and the obtained orbitals and local basis is complemented with a monocentric Gaussian basis and B-splines.

In this way, the Xchem code is divided in modules with different performance: - Module basis: Prepares all 1 electron integrals and it is not parallelized. - Module orbitals: Creates a set of linear independent orbitals with the monocentric basis and orthogonalized them to the local orbitals. It is parallelized at OpenMP level. -Module integrals: This module is one of the most critical ones and it is barely parallelized at OpenMP level. However, the different calculations can be easily parallelized by sending them to different nodes. For a standard calculation of 22 monocentric Gaussian exponents, the module can be parallelized up to 253 nodes (usually we used 10). -Module rdm: Calculates the Reduced Density Matrices of the different electronic states and can be highly parallelized by sending them to different nodes. Moreover, this calculation is prepared to be independent of the integrals one, sending them to different nodes.

E. Performance/Scaling of the ePolyScat Code

Eliminate ePolyScat entirely if this information is not provided.

F. Performance/Scaling of the UKRMol Code

The UKRmol+ code is parallelized using a combination of OpenMP and MPI. Table I shows the scaling of the UKRmol+ code on ARCHER for the diagonalization of the Hamiltonian matrix in an electron-uracil calculation. The “Expected” scaling is given approximately by $(\text{Configurations})^3/(\text{tasks})$.

| MPI tasks | Configurations | Time | Expected |
|-----------|----------------|---------|----------|
| 32 | 19668 | 10m 48s | - |
| 121 | 39725 | 19m 41s | 23m 32s |
| 2025 | 135835 | 37m 39s | 56m 13s |

TABLE I. UKRmol+ Hamiltonian diagonalization for electron-uracil calculation for the number of configurations indicated run on ARCHER. The ‘Expected’ scaling is given approximately by $(\text{Configurations})^3/(\text{tasks})$

Figure 4 shows the scaling of the construction of the Hamiltonian matrix by MPI-SCATCI for phosphoric acid. The linear scaling begins to deteriorate at about 32 cores.

122815 X 122815 Hamiltonian build for phosphoric acid

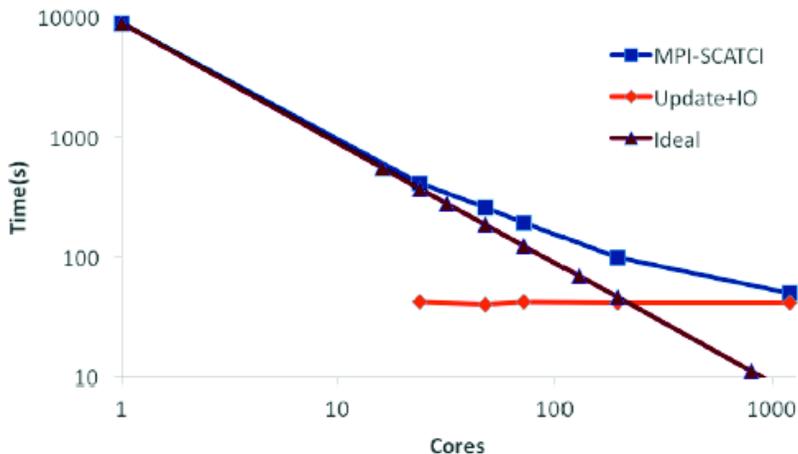


FIG. 4. UKRmol+ Hamiltonian construction for electron-scattering from phosphoric acid using the program MPI-SCATCI. The update+IO time is the time taken for MPI synchronization steps that include disk writes. The time taken by MPI-SCATCI includes the update+IO time.

G. Performance/Scaling of the RMT Code

RMT uses the R -matrix paradigm, partitioning the interaction region into an “inner” and an “outer” region with different numerical schemes utilized in each region. This approach facilitates efficient parallelization without sacrificing accuracy. Figures 5–7 provide some information about the strong scaling of an entire calculations, as well as individually in the inner and outer regions, respectively.

Strong scaling cannot be trivially inferred from one calculation, as at some point the bottleneck switches between the two regions (inner/outer). It is relatively easy to demonstrate scaling if one region dominates, but the more difficult problem in setting up efficient calculations is finding the right balance between the two regions. Both regions need to be optimized in order to see full scaling, and the optimal arrangement is likely to be quite different for the two regions, and quite different in each specific calculation.

Limited RMT scaling data are available on XSEDE machines. Most of the data below come from the ARCHER (Cray XC-30). Moreover, the most recent capabilities and performance enhancements in RMT have not yet been tested for scaling on any machine. These are both reasons for inclusion in the current proposal.

Table II displays the weak scaling behaviour of the finite-difference based RMT outer region. The number of cores used in the outer region determines the total size of the configuration space. Increasing the number of outer region cores has a negligible effect on the average iteration time (i.e. the time taken to propagate the wave function one step forward in time). Thus efficient scaling of the outer region is possible to very large distances without increasing overheads.

| Cores | 620 | 750 | 1000 | 1500 | 2000 | 2500 |
|--------------------|-------|------|-------|-------|-------|-------|
| Iteration Time (s) | 1.135 | 1.08 | 1.134 | 1.135 | 1.135 | 1.148 |

TABLE II. Weak scaling of outer region: Argon, two target states, inner region fixed (144 cores) on ARCHER (CRAY XC30).

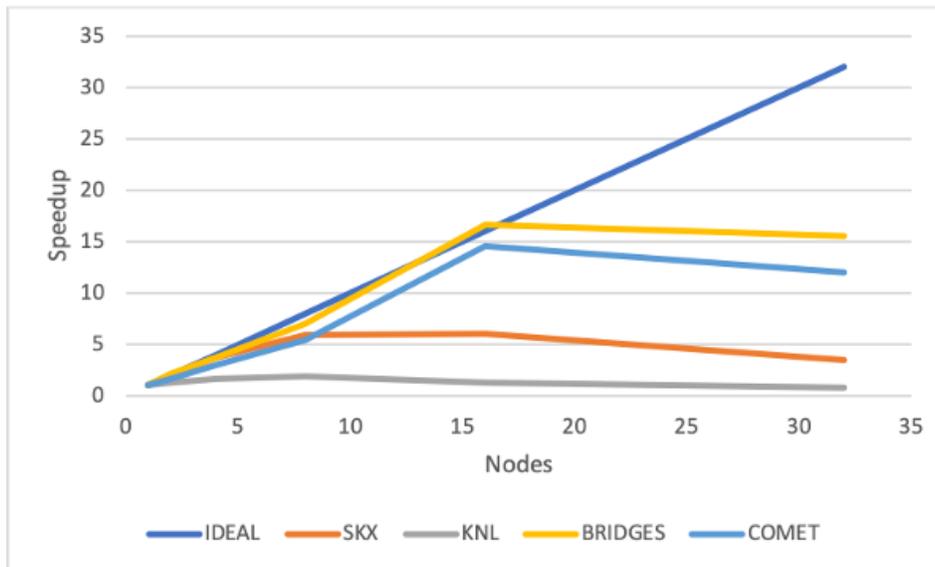


FIG. 5. Strong scaling of RMT in an entire calculation: Photoionization of Argon with circularly polarized light (XSEDE machines). Preliminary scaling data have been obtained for RMT on several XSEDE machines (Comet, Bridges and Stampede2). In this calculation the ratio between the number of inner region cores and outer region cores was kept fixed. Calculations were performed with the same level of optimization on all machines, although the number of cores per node varies from 48 cores/node on Stampede2 (SKX and KNL) to 28 and 24 cores/node on Bridges and Comet, respectively.

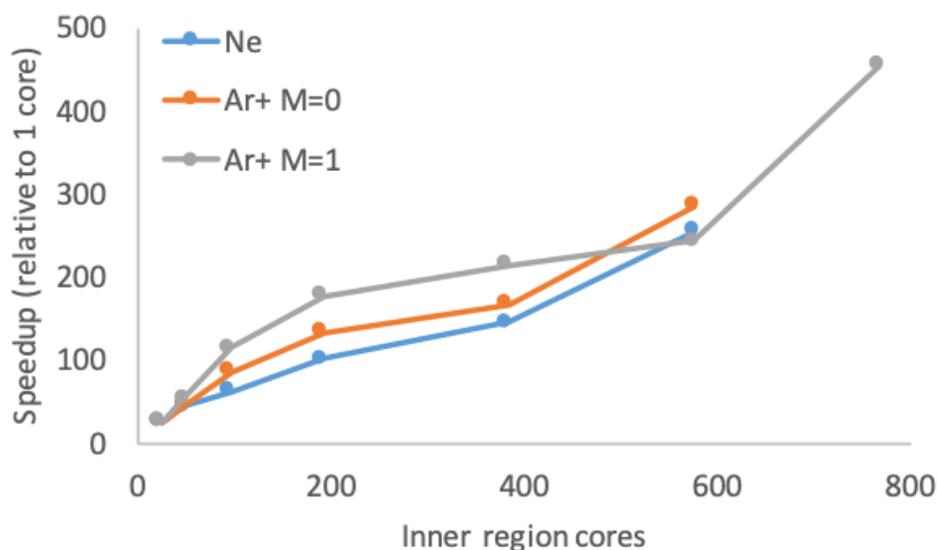


FIG. 6. Strong scaling of RMT in the inner region: Various target atoms, outer region fixed (1200 cores) on ARCHER (CRAY XC30). The more complete the description of atomic structure, the larger the matrix-vector multiplication in the inner region. Calculations using non-linear laser polarisation will be even more intensive than the largest (Ar+ M=1) calculation shown here, and hence we expect that higher core counts will be possible.

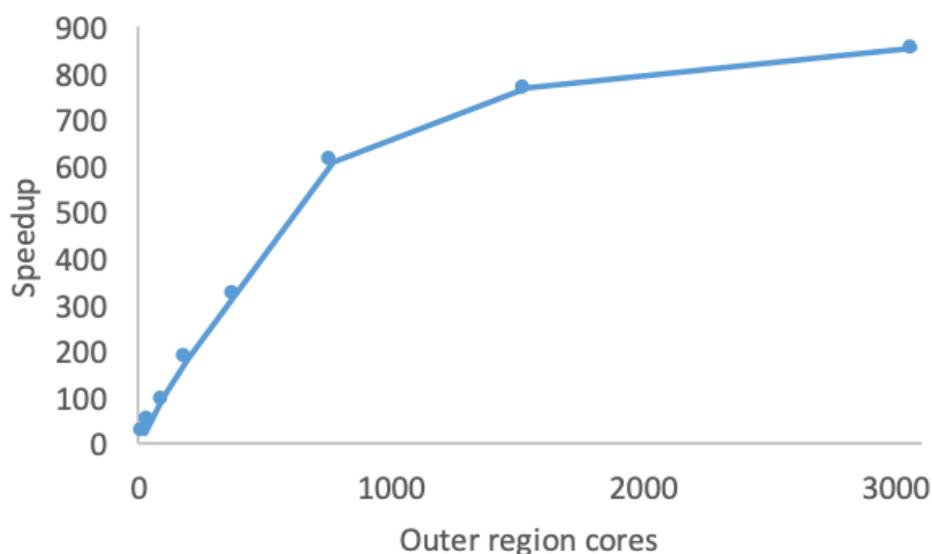


FIG. 7. Strong scaling of RMT in the outer region: Argon, two target states, inner region fixed (144 cores) on ARCHER (CRAY XC30). For calculations comprising many electron emission channels, the outer region becomes the bottleneck. Performance is enhanced by reducing the size of the radial grid on each core. The levelling off above 1,000 cores shows that the inner region is becoming dominant. Additional performance could then be gained by adding cores to the inner region.

References

- [1] <http://www.netlib.org/scalapack/>
- [2] J. Zhu and A. Scrinzi, arxiv.org/abs/1912.09250.