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. Fig. 1 provides some information about the strong scaling of an entire calculation, while Fig. 2 shows the scaling for inner and outer regions, side-by-side.

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 different for the two regions, and again 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 I 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 needed 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 1: Weak scaling of outer region: Argon, two target states, inner region fixed (144 cores) on ARCHER (CRAY XC30)

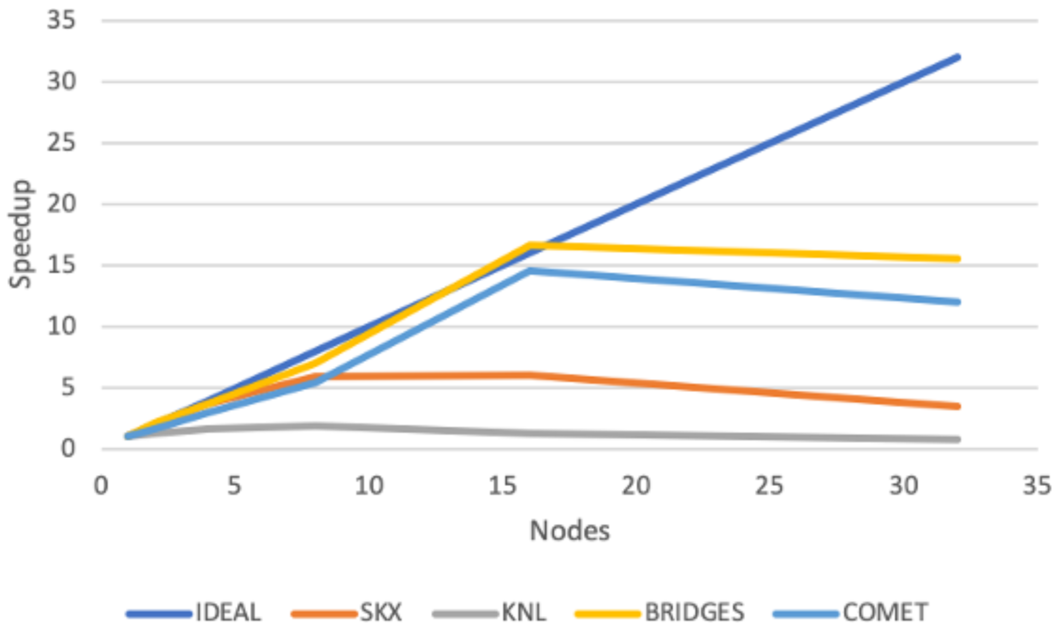

FIG. 1. Strong scaling of RMT in an entire calculation: Photoionization of Argon with circularly polarized light (XSEDE machines). Preliminary scaling data have been obtained for RMT on several XSEDE machines (Comet, Bridges and Stampede2). In this calculation the ratio between the number of inner region 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.

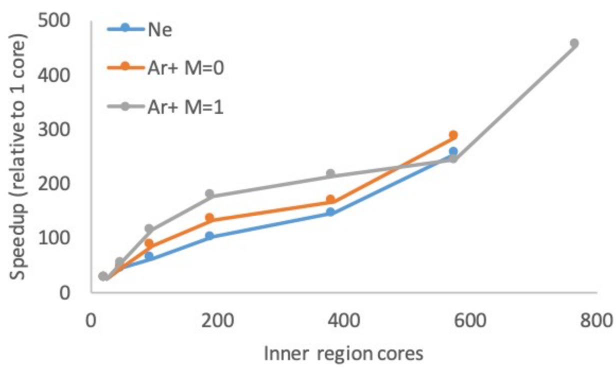

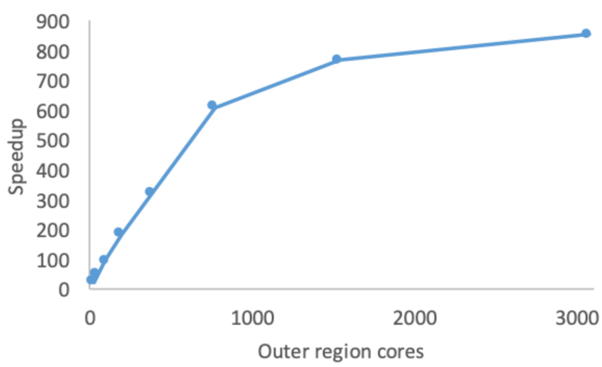

FIG. 2 Left. Strong scaling of RMT in the inner region: Various target atoms, outer region fixed (1200 cores) on ARCHER (CRAY XC30). The more complete the description of atomic structure, the larger the matrix-vector multiplication in the inner region. Calculations using non-linear laser polarisation 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. 2. Right. 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.