


Performance as a function of the number of GPUs. Dashed red line is ideal scaling, blue shows particle time, black shows total time. Degradation of total time is due to the all to all transpose in the FFT, which dominates. The FFT takes about 40-70% of the total time. Single precision with Cuda Fortran



Performance as a function of the number of GPUs. Dashed red line is ideal scaling, blue shows particle time, black shows total time. Degradation of total time is due to the all to all transpose in the FFT, which dominates. The FFT takes about 30-45% of the total time. Double precision with Cuda Fortran.