TRINITY

SAND2017-9703R

Scaling tests of a new algorithm for DFT hybrid-functional calculations on Trinity Haswell

Alan Wright and Normand Modine

Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Los Alamos NATIONAL LABORATORY EST. 1943

Sandia National Laboratories

NNSA National Nuclear Security Administration

Alliance for Computing at Extreme Scale

Objective and Approach

Objective: Develop a highly scalable implementation of Density-Functional-Theory (DFT) with hybrid functionals (combining density- and orbital-dependent electron exchange) that can be used in Sandia *Radiation-Effects* and *High-Energy-Density-Physics* Mission Areas

Approach:

- Used Sandia's Socorro DFT code [1] as a development platform
- Implemented a Compact Representation (CR) [2] of the orbital-dependent exchange operator, which is constructed from derivatives of the exchange energy with respect to the orbitals…

$$\frac{\delta E_x}{\delta \psi^*_{n_1\vec{k}_1}(\vec{r}_1)} = f_{n_1\vec{k}_1} \sum_{\vec{k}_2} \sum_{n_2} \left(-\frac{\Omega}{2} e^2 f_{n_2\vec{k}_2} \psi_{n_2\vec{k}_2}(\vec{r}_1) \int_{\Omega} \frac{\psi^*_{n_2\vec{k}_2}(\vec{r}_2)\psi_{n_1\vec{k}_1}(\vec{r}_2)}{|\vec{r}_1 - \vec{r}_2|} d\vec{r}_2 \right)$$

- Developed a new algorithm that allows use of integer multiples (> 1) of $N_{orbitals}$ x $N_{k\text{-}points}$ MPI processes when constructing the CR, and multiple threads for each MPI process
- Tested and used the new algorithm under *Open Science 1* and *ATCC 1-3*

[1] dft.sandia.gov/socorro
[2] Duchemin and Gygi, Comput. Phys. Commun. **181**, 855 (2010)

A Simple Example of the New Algorithm

5 orbitals and 10 MPI processes

Distribution of the LHS (row, R#) orbitals

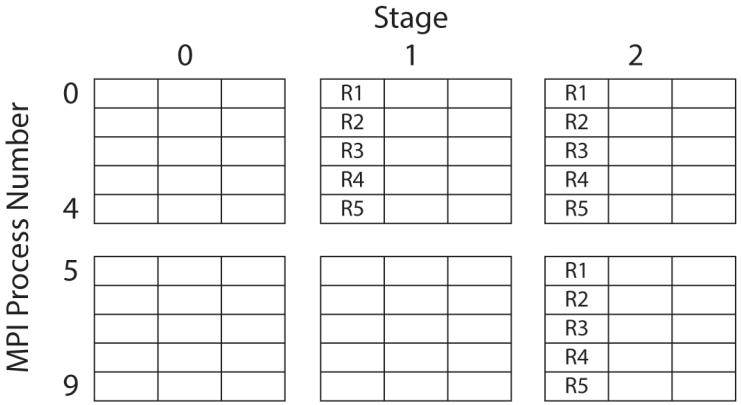

Distribution of the RHS (column, C#) orbitals

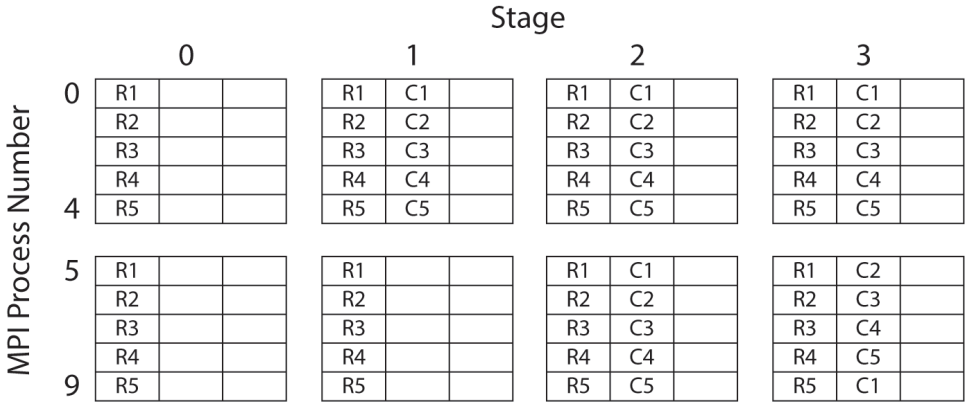

Accumulation of the derivatives wrt the row orbitals (DR#)

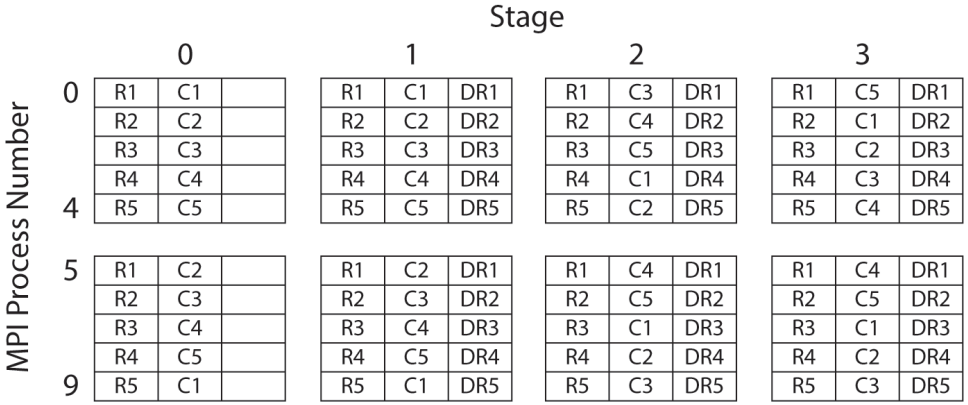

Collection of the derivative contributions

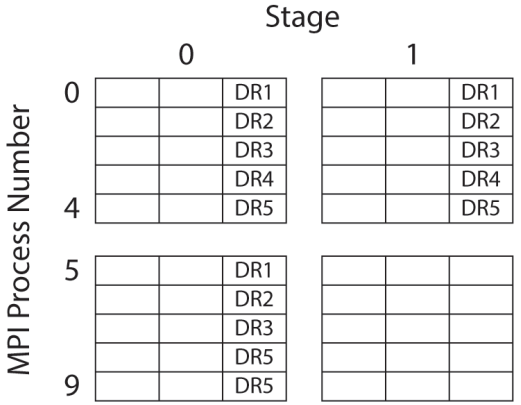

Scaling Tests

Overview We show scaling results for materials of interest in Sandia *Radiation-Effects* and *High-Energy-Density-Physics* Mission Areas. Each timing is from a self-consistent calculation for bulk material. Two timings are given: (1) walltime for the construction of the CR exchange operator (*Exchange-Operator*) and (2) walltime for everything else (*non-Exchange-Operator*).

GaAs 216-atom cell: We show timings from runs using 8 k-points, 640 orbitals, a plane-wave cutoff of 40 Ryd to expand the orbitals, a PBE(23) functional [3] and up to 286,820 CPUs.

Gold: We show timings from runs in a 256-atom cell using 1 k-point, 2304 orbitals, a plane-wave cutoff of 50 Ryd to expand the orbitals, the HSE functional [4], and up to 73,720 CPUs.

GaAs 512-atom cell: We show timings from runs using 8 k-points, 1280 orbitals, a plane-wave cutoff of 40 Ryd to expand the orbitals, a PBE(23) functional and up to 286,820 CPUs.

[3] J. Chem. Phys. **110**, 1029 (1999); Phys. Rev. Lett. **77**, 3865 (1996)
[4] J. Chem. Phys. **118**, 8207 (203)

Timings for GaAs in a 216-atom cell w/o threads:

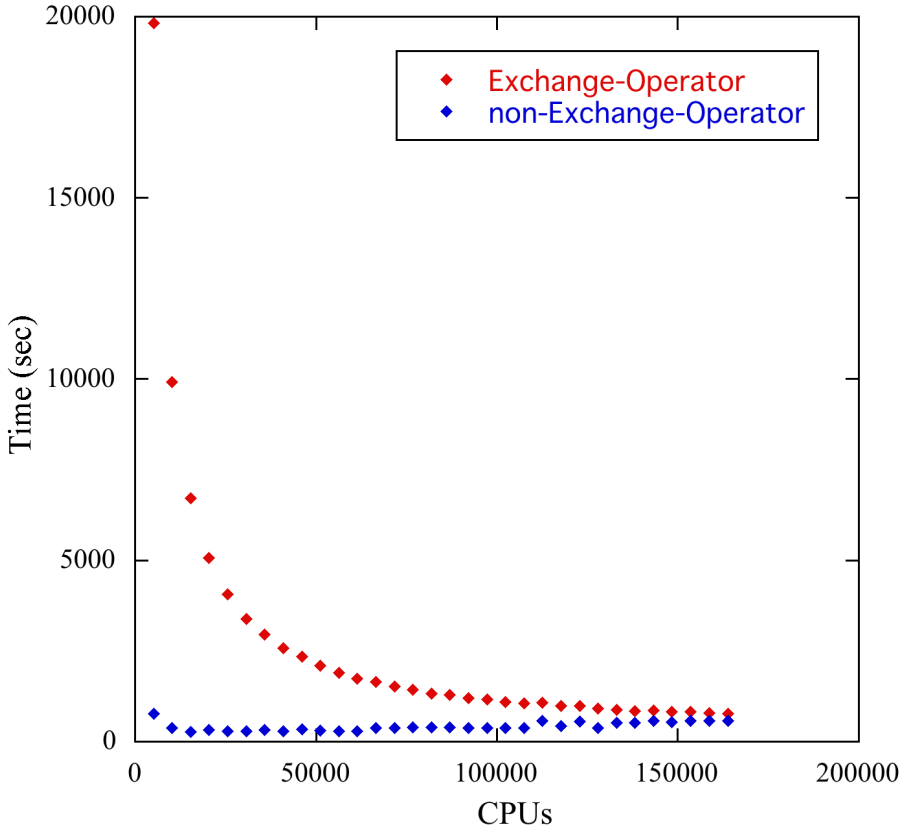

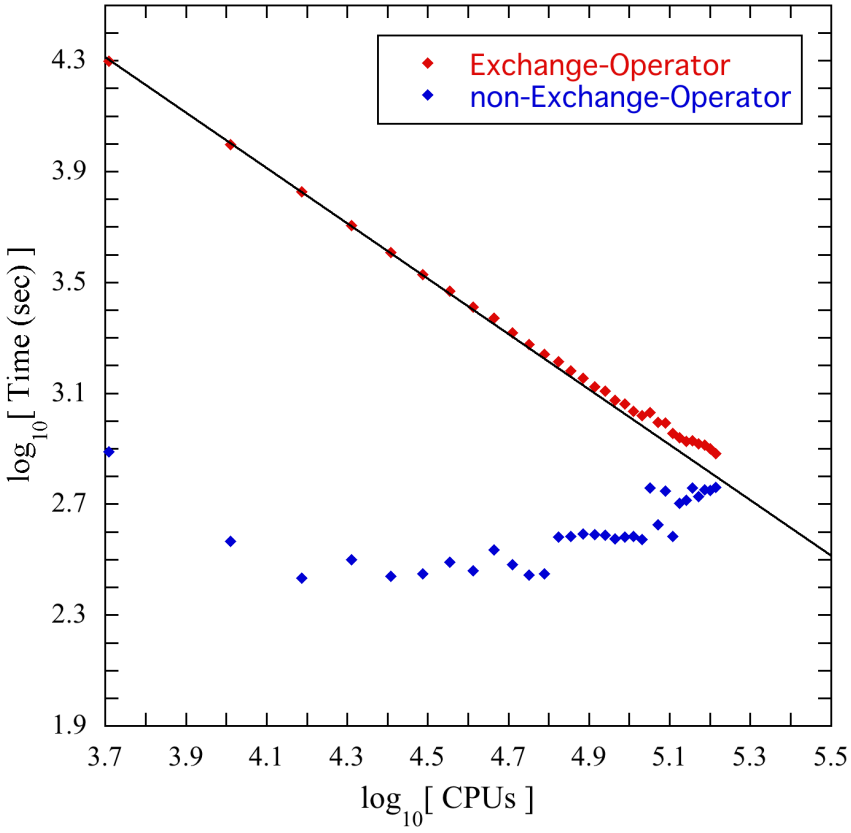

The line in the log/log plot gives the slope one should find with perfect scaling - hereafter we only show log/log plots in order to clearly reveal the scaling

Timings for GaAs in a 216-atom cell w/ and w/o threads:

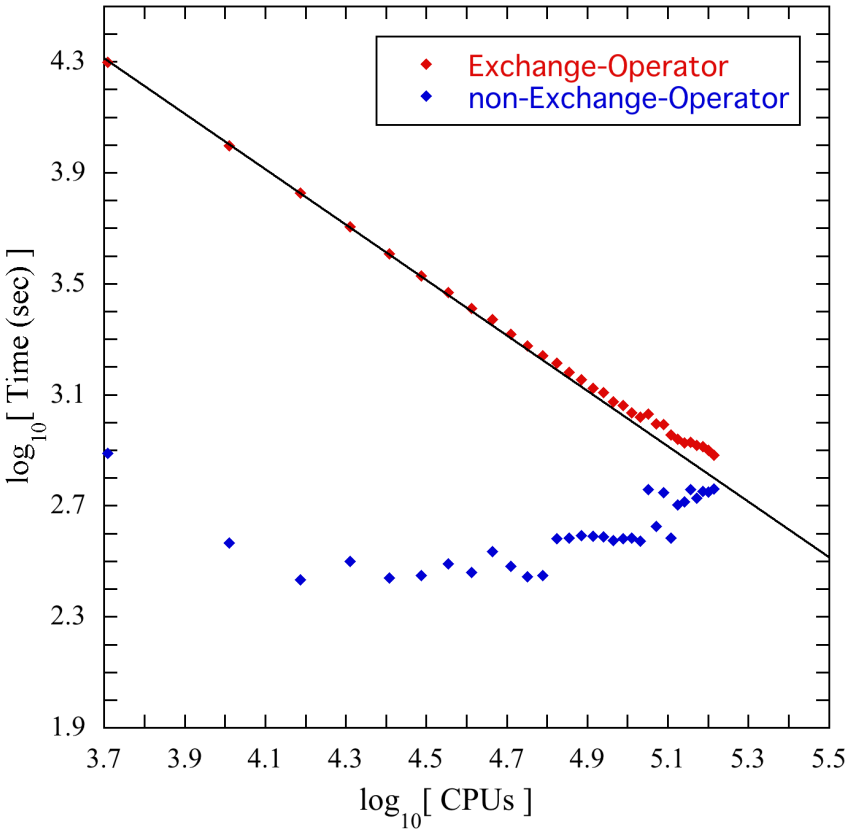

No threads

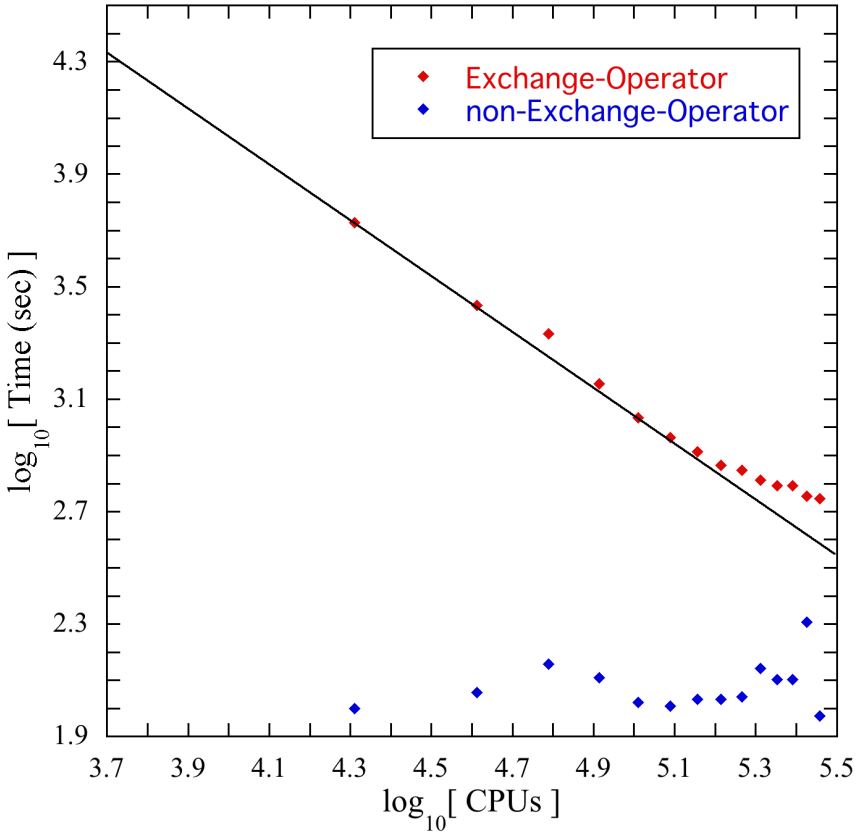

4 threads per MPI process:
The scaling range is extended and
non-Exchange-Operator times are reduced

Timings for gold with w/ and w/o threads:

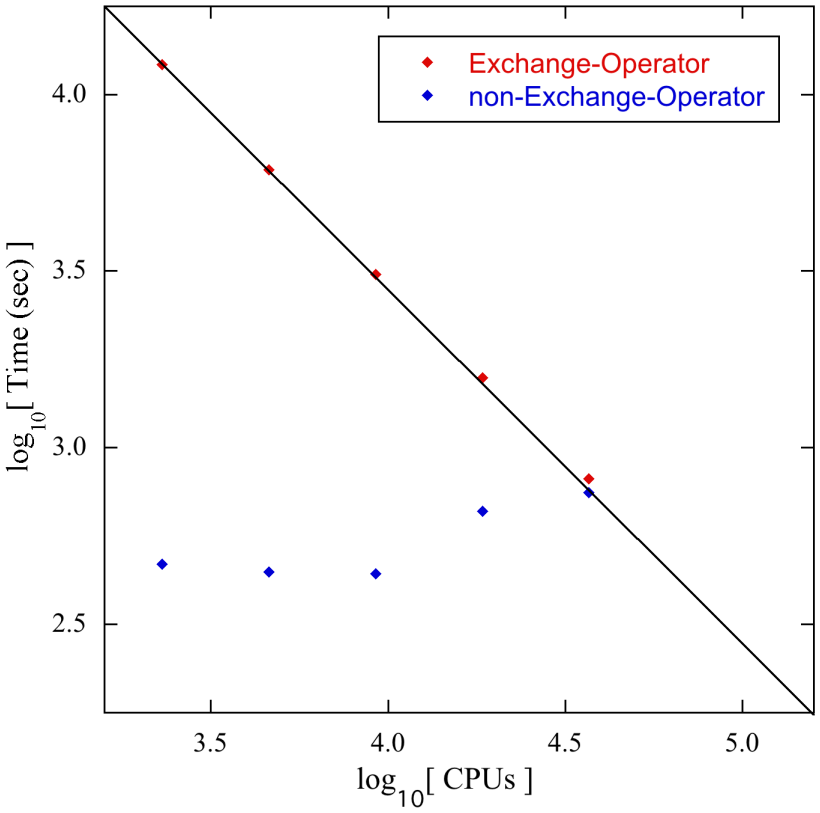

No threads

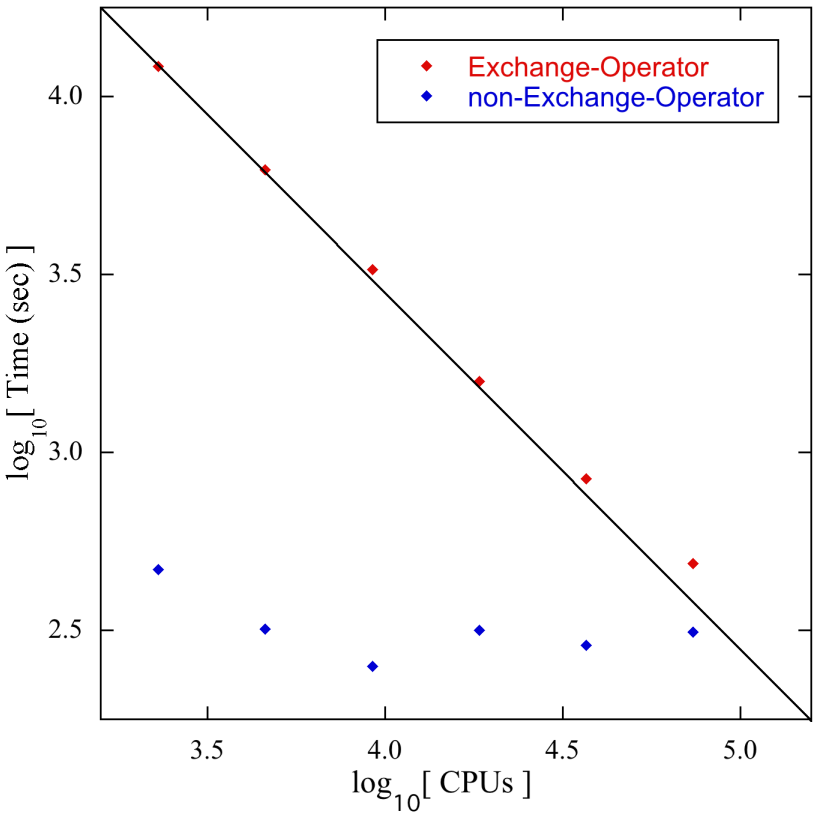

Multiple threads per MPI process"
Scaling is extended and
non-Exchange-Operator times are reduced

Timings for GaAs in a 512-atom cell w/ threads:

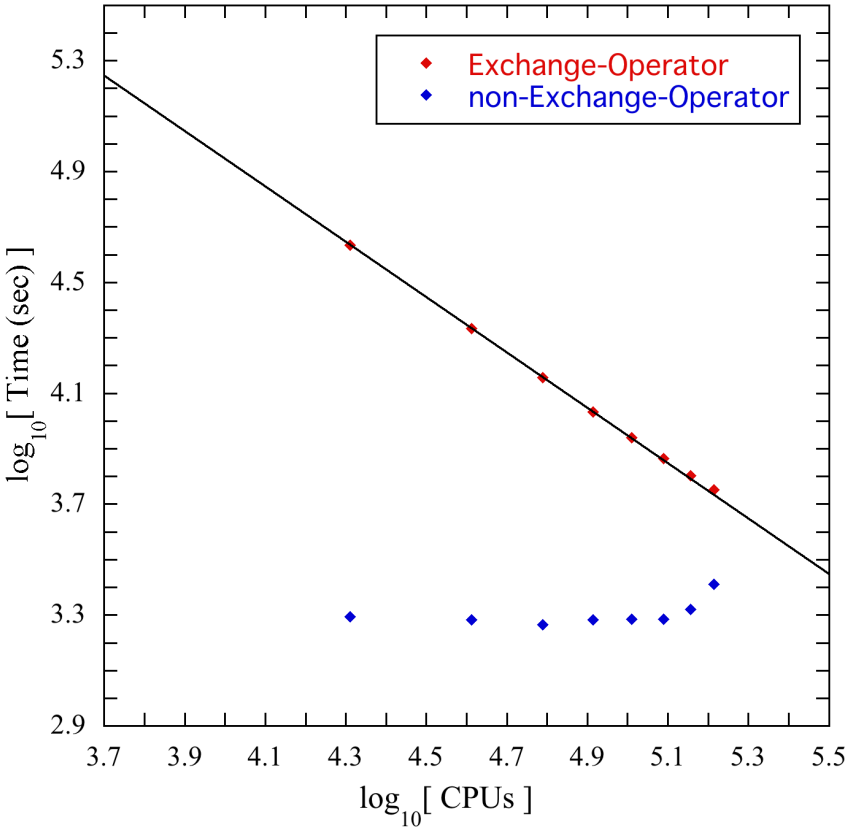

2 threads per MPI process

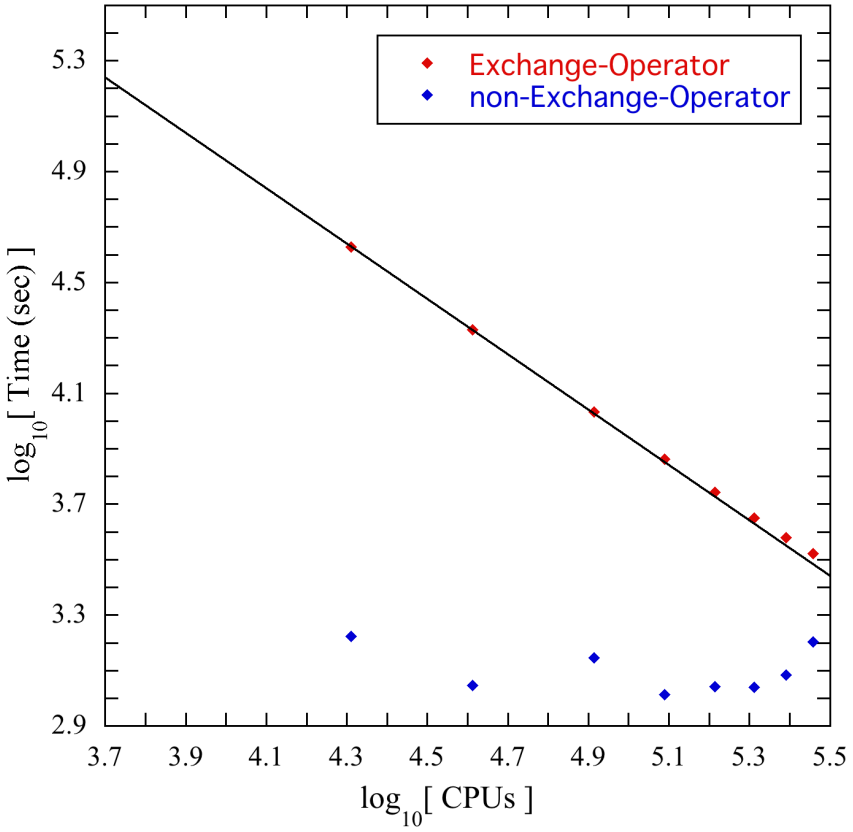

4 threads per MPI process:
Scaling runs limited by machine size

Our Assessment of Trinity

Hardware: We find the Trinity (Haswell) hardware to be fast and reliable, and have used Trinity to study migration processes of a point defect in GaAs

User Support:

- With the help of Mike Davis and Mahesh Rajan we were able to quickly get Socorro running on Trinity with excellent performance
- We requested and were given Dedicated Application Time for some of our early scaling tests
- With the help of Joel Stevenson we easily made the transition from Moab to Slurm
- In every aspect, the level of user support we received was excellent

Overall: We find Trinity to be ready for production work in the Sandia Mission Areas noted in the introduction

Future Needs: We expect that we could efficiently use a Trinity-type machine several times larger than Trinity for anticipated work in the *Radiation-Effects* Mission Area

Our Sincere Thanks to the High-Performance Computing Staffs at SNL and LANL, and to ASC