High Performance Computing Workshop

DOE CSGF Fellow and Alumni HPC Stories: Jack Deslippe

Research Goals

- **Predict the structural, electronic and optical properties of any material from first-principles.**

- **Guide experimentalists in the search for novel materials and physical understanding of phenomena in complex systems.**

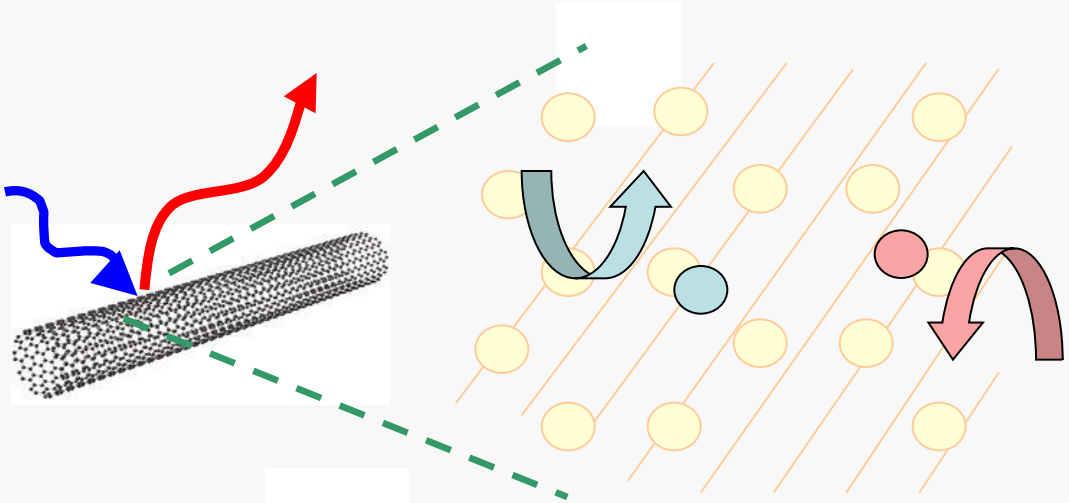

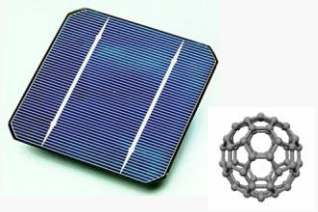

Nano Photovoltaics

Easy Right?

- **"If you can't do Silicon, you can't do anything."**
- **"If you can do Silicon, it doesn't mean you can do anything else."**
- **For excited electronic properties we use the "GW-Bethe-Salpeter Equation" methodology.**
- **It is an involved process with lots of Linear Algebra FFTs, inversion, diagonalization**

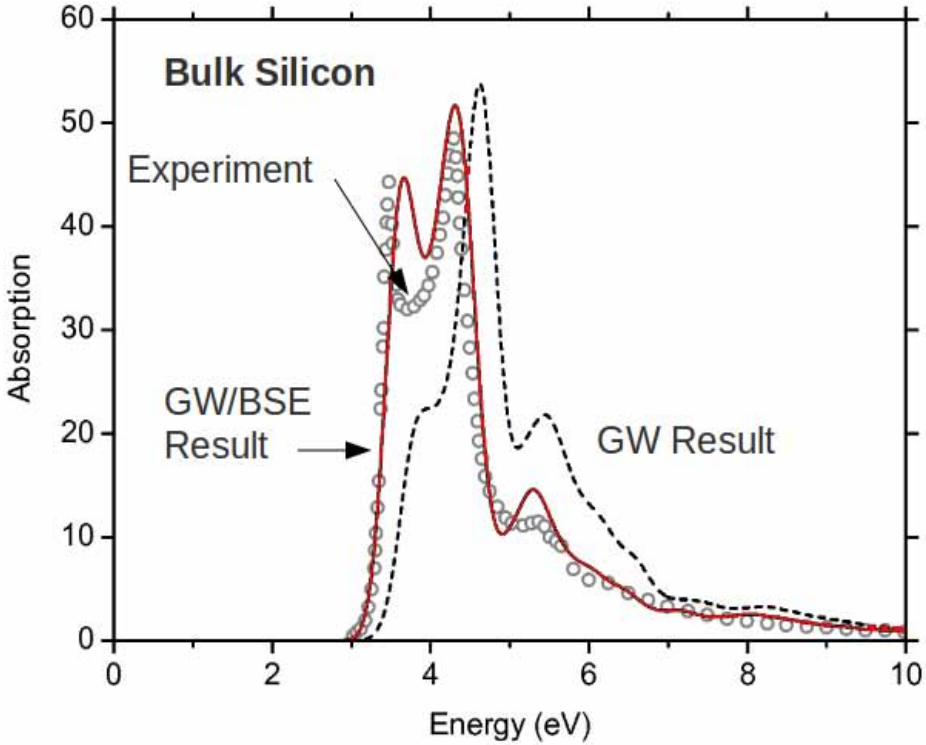

Expt. G.E. Jellison, M.F. Chisholm, S.M. Gorbatkin, Appl. Phys. Lett. 62, 3348 (1993).

Not So Fast

- **In 2006 - studying large nano-systems (~100s of atoms) = Failure.**

PROBLEMS:

- **Code optimized/parallelized for small systems. Max of ~100 processors.**
- **Communication done through disk**
- **Significant memory limitations. Essentially no distribution among processors.**
- **15 different versions of the same code with no organized development.**

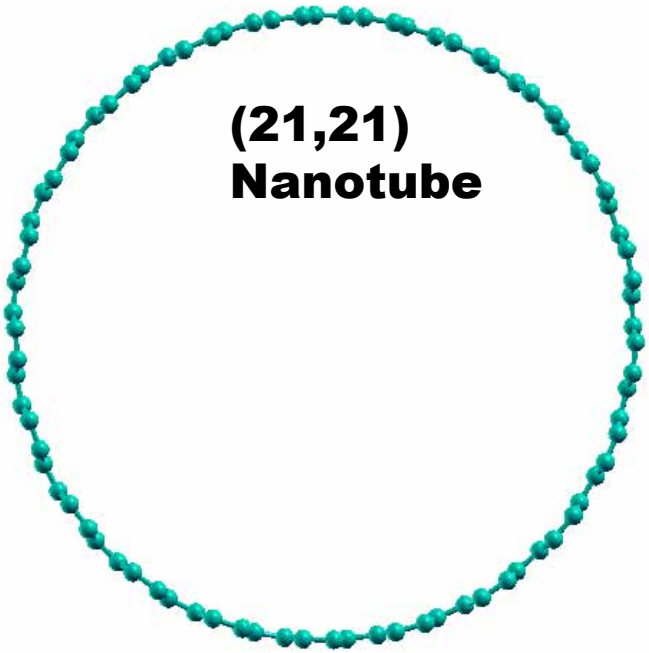

Wavefunctions – 100GB
Dielectric Function – 120 GB
Exciton Hamiltonian – 200 GB

Overcoming Obstacles

Step 0 – Pick a logo.

Step 1 - Cleaning up the mess. Introduced the Group to SVN (Subversion) and Trac. Merged the 15 versions into 1.

3 years later ~ 2000 commits in the SVN repository from over 20 students/postdocs/visitors. Robust, centralized version of the code.

Step 2 – Parallelization: Change communication through disk to MPI with nearly perfect linear scaling of memory.

Step 3 – Optimize the bottlenecks. Find the slowest part of the code and optimize it for thousands of processors.

Example Barrier

Naïve scaling – NG^4
Actual scaling – NG^3 ln(NG)
Parallel scaling – NG x ln (NG) (parallel over *n*,*n*')

Level 3 Blas – (1-5% of calculation for largest systems)

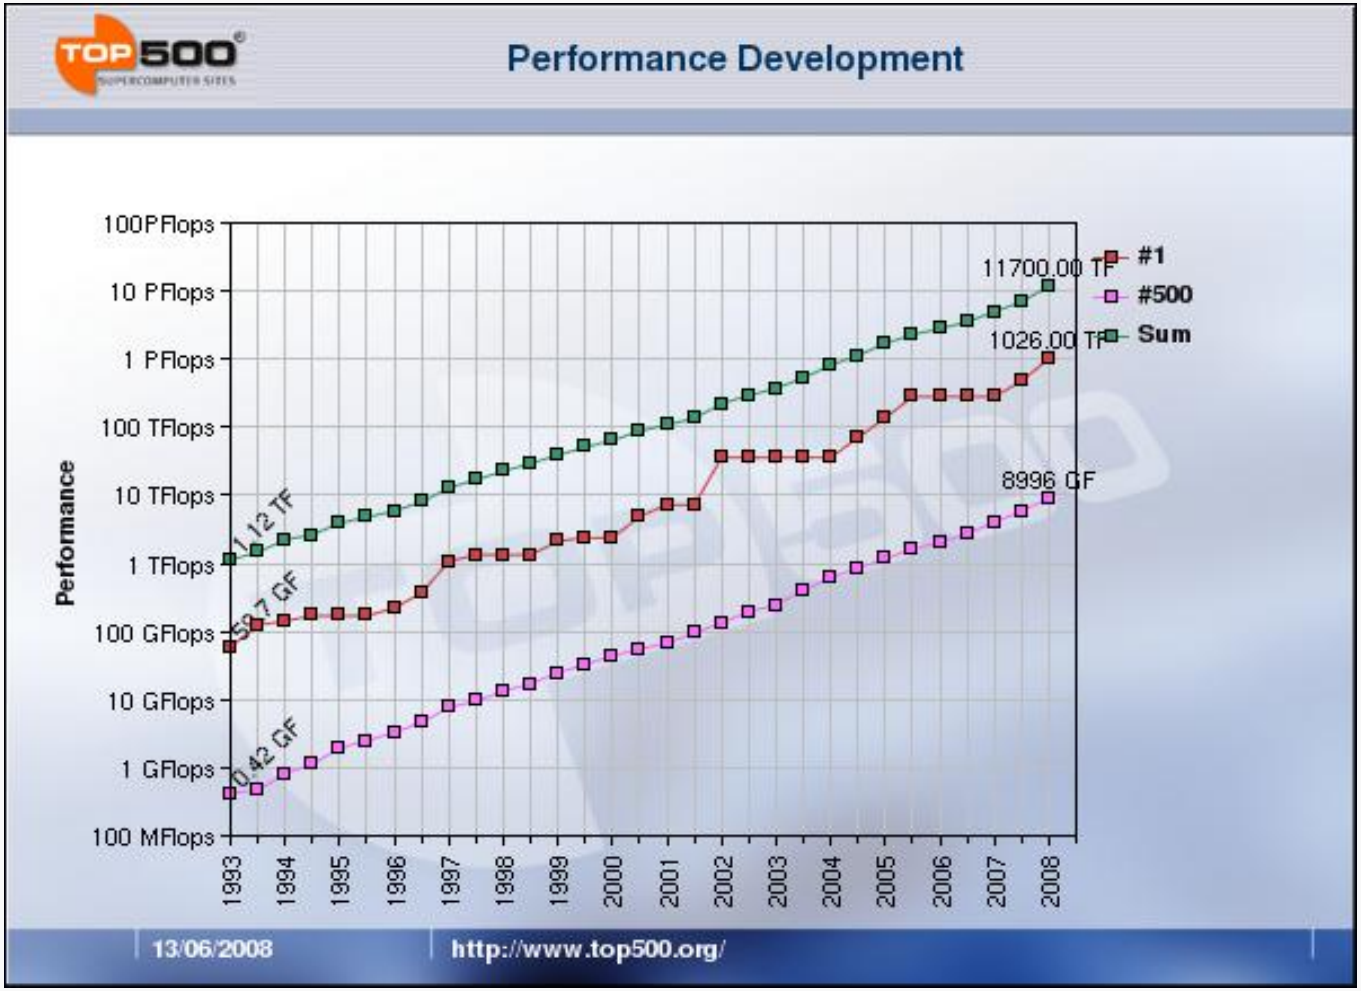
TOP500
SUPERCOMPUTER SITES
Performance Development
100 PFlops
10 PFlops
1 PFlops
100 TFlops
10 TFlops
1 TFlops
100 GFlops
10 GFlops
1 GFlops
100 MFlops
Performance
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
#1
#500
Sum
11700.00 TF
1026.00 TF
8996 GF
1.12 TF
59.7 GF
0.42 GF
13/06/2008
http://www.top500.org/

Avg CPUs on Machine on Top 500

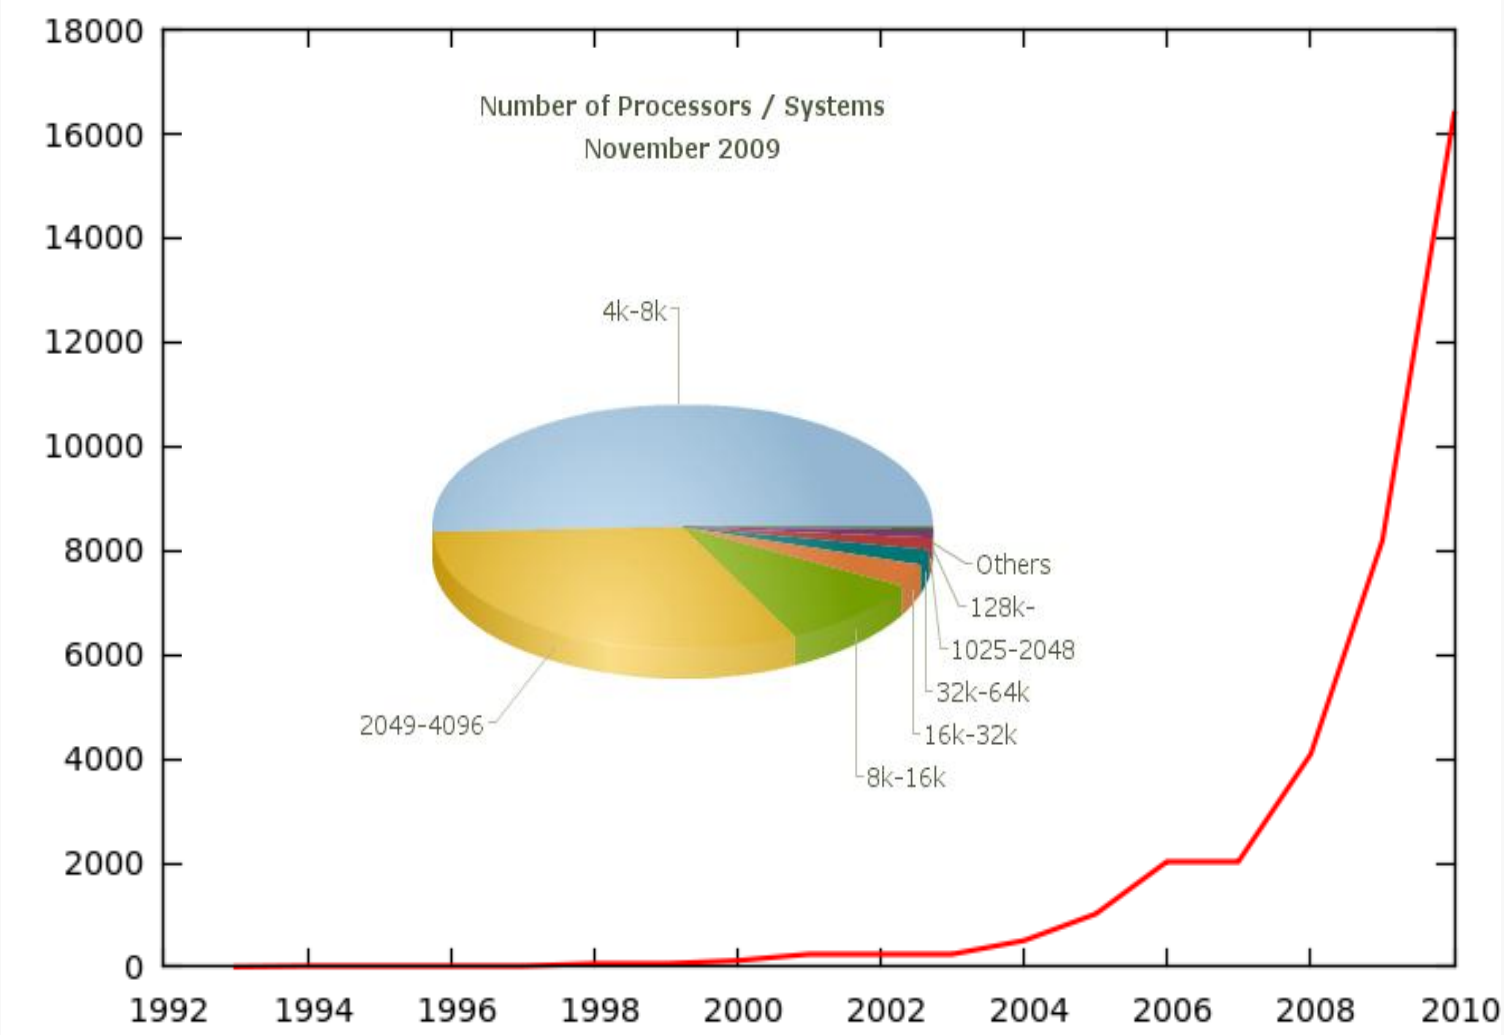

The Pay Off

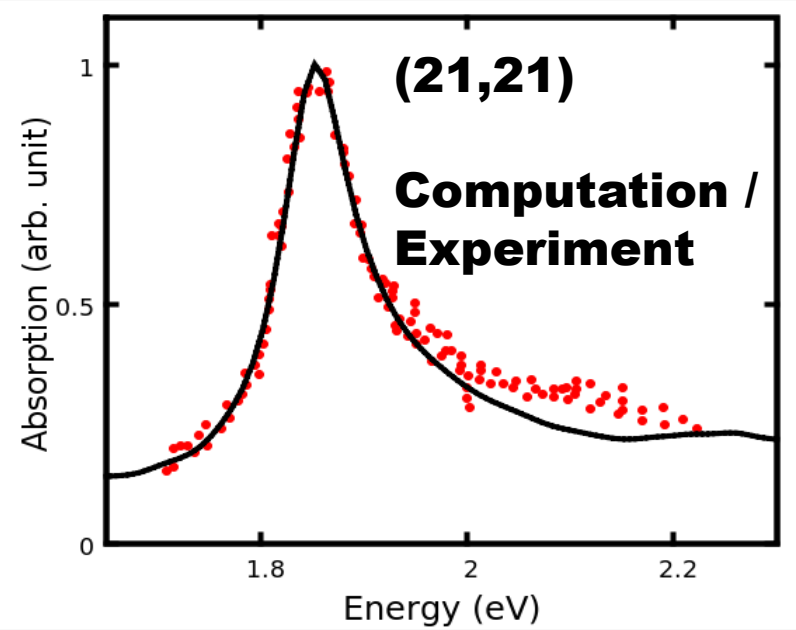

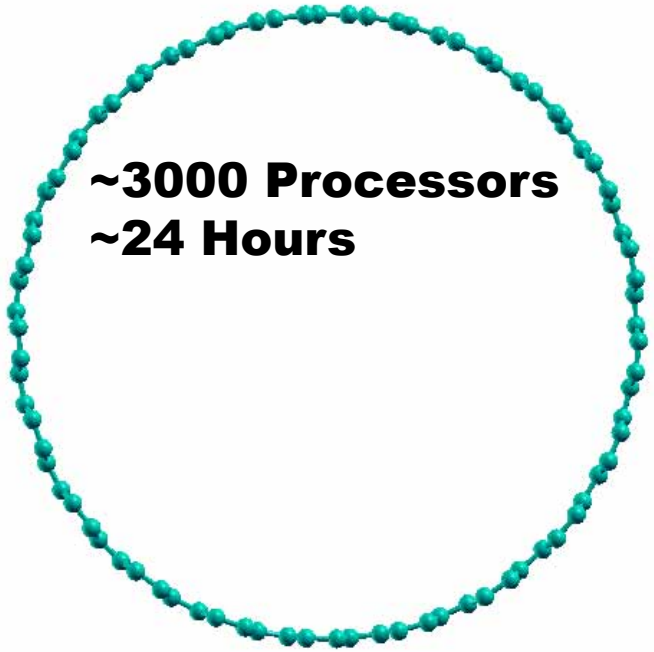

Observation of Excitons in One-Dimensional Metallic Single-Walled Carbon Nanotubes

Feng Wang,[1] David J. Cho,[1] Brian Kessler,[1] Jack Deslippe,[1] P. James Schuck,[2] Steven G. Louie,[1,2,3] Alex Zettl,[1,3] Tony F. Heinz,[4] and Y. Ron Shen[1,3]

Scaling

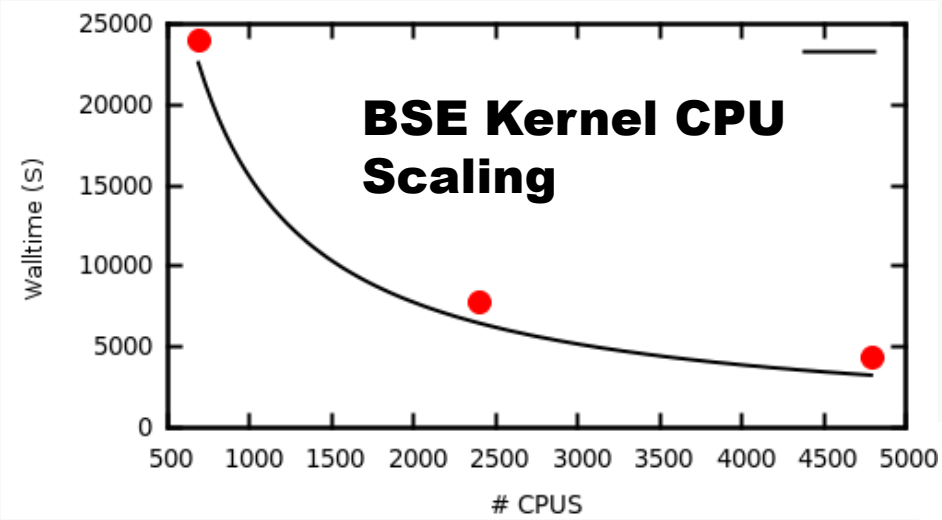

(21,21) Nanotube

Equivalent to ~500 Bulk Atoms

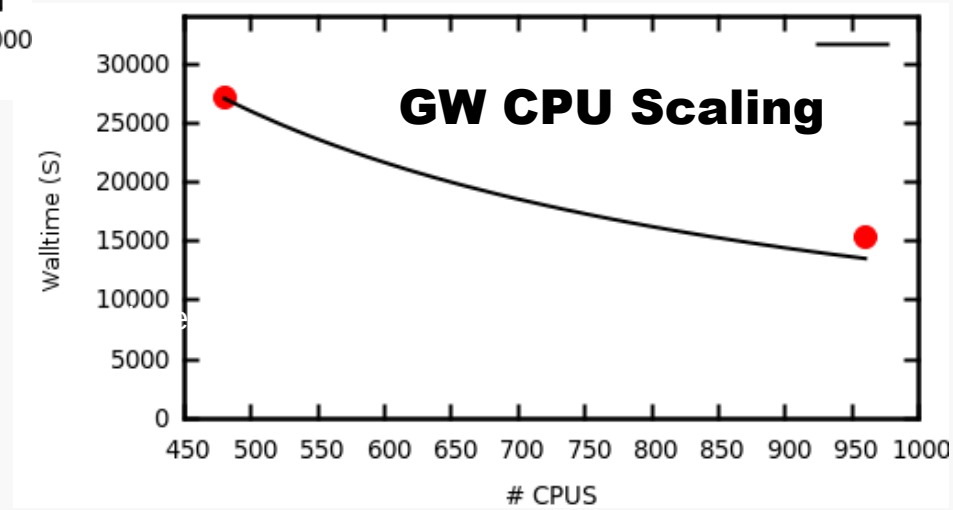

Lessons Learned

- **Utilize development tools. Professional developers use them for a reason.**
- **Running on thousands of CPUs is much better than running on tens/hundreds.**
- **Don't reinvent the wheel – utilize libraries BLAS/LAPACK/scaLAPACK etc… whenever possible.**
- **Parallelize both for time and memory.**