

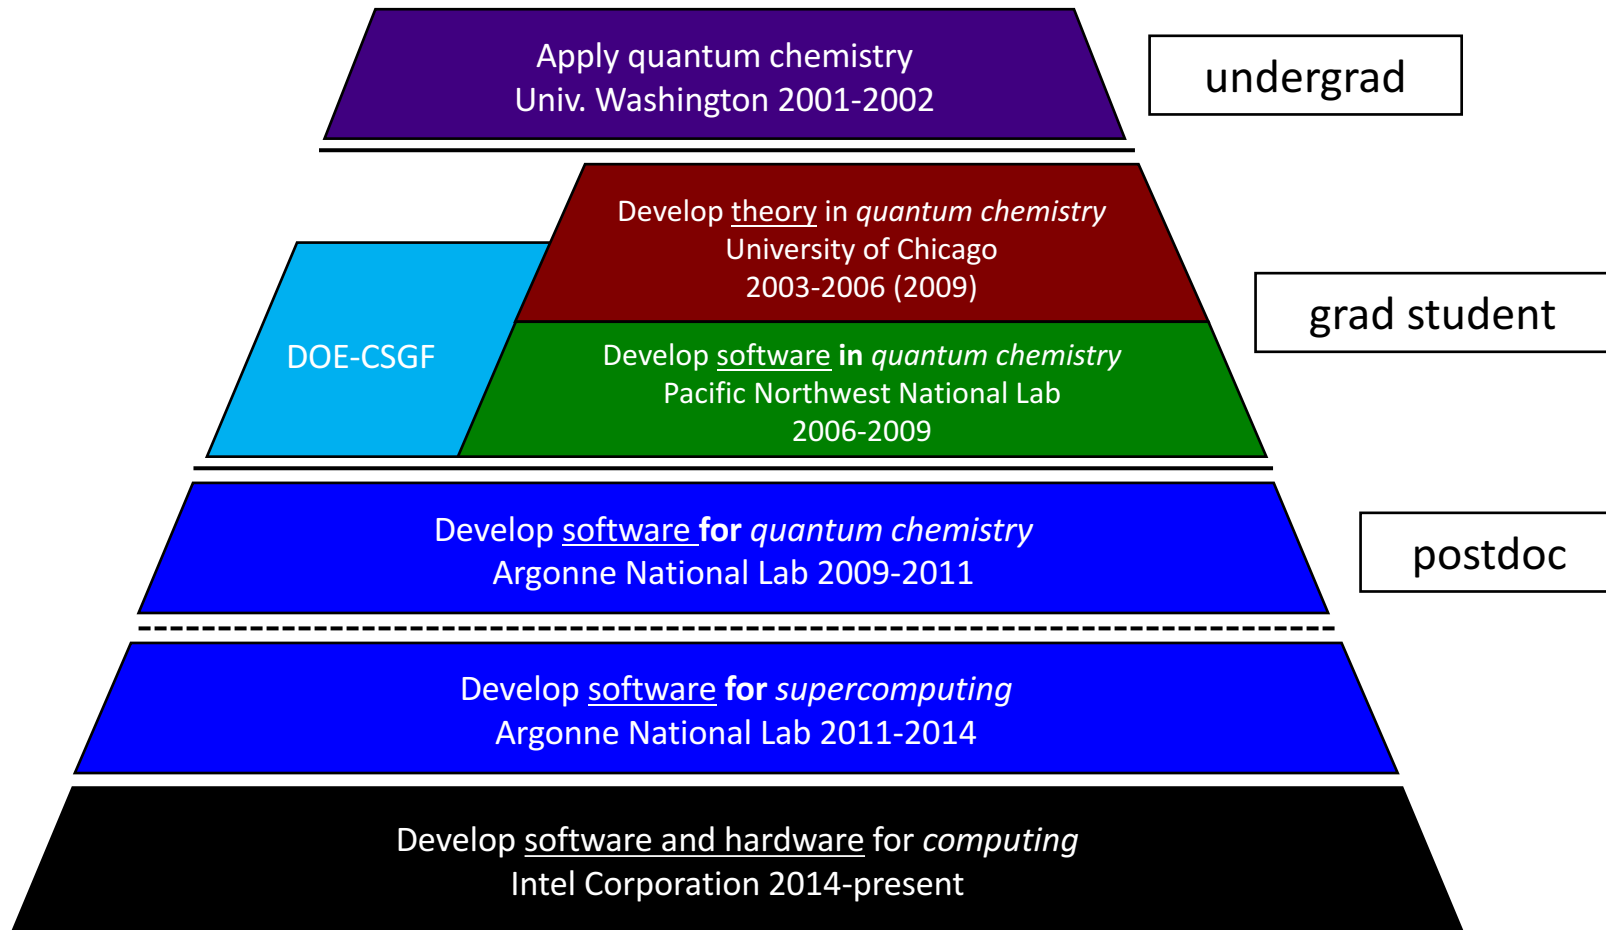
Down the rabbit hole: From B3LYP to x86

Jeff Hammond
DOE-CSGF 2005-2009

jeff.science@gmail.com
<https://jeffhammond.github.io/>

Disclaimer

- I am speaking in a personal capacity as a DOE-CSGF alumnus, not as an Intel employee.
- My views do not represent those of:
 - Krell Institute
 - Argonne National Lab
 - Pacific Northwest National Lab
 - Anybody else
- Hanlon's Razor (blame stupidity, not malice).



Apply quantum chemistry
Univ. Washington 2001-2002

undergrad

Develop theory in *quantum chemistry*
University of Chicago
2003-2006 (2009)

grad student

DOE-CSGF

Develop software in *quantum chemistry*
Pacific Northwest National Lab
2006-2009

Develop software **for** *quantum chemistry*
Argonne National Lab 2009-2011

postdoc

Develop software **for** *supercomputing*
Argonne National Lab 2011-2014

Develop software and hardware for *computing*
Intel Corporation 2014-present

B.C. (Before CSGF)

- Undergrad at University of Washington.
 - BS Chemistry, BA Math.
 - Research in computational chemistry applied to organic chemistry.
 - No computer science courses (discouraged by chemists: “why learn Java?”)
 - Ran Gaussian and Molcas on DEC Alpha and IBM RS6000 workstations.
- Early grad school
 - Pencil-and-paper theory augmented with Matlab prototypes.
 - Coursework diversity in mathematical physics, quantum computing (CS dept)

Throughout this period, I wanted to move into HPC but had no idea how to get there.

Discovering CSGF

- When I found the CSGF program, I was immediately drawn to it:
 - POS goals were consistent with my interdisciplinary interests.
 - Provided an immediate connection to the HPC community that I did not have.
- If at first you don't succeed...
 - My first application to CSGF as G1 was not successful.
 - Received the fellowship at G2.
 - Eligibility changes no longer allow G2 applications...

Program of Study

- ...was not always fun: do homework or write Physical Review paper?
- Things I never would have learned otherwise:
 - Random number generators
 - Parallel computer architecture (w/ Rick Stevens)
 - Write PDE solvers in MPI, ran on Blue Gene/L (w/ Andrew Siegel)
 - The difference between programming and computer science 😊
 - Computer scientists do not speak Fortran
- I failed (to get a B in) multiple courses...
 - Algorithms (twice)
 - Numerical methods in PDEs

The Practicum

- I chose PNNL because:
 - The DOE program manager in comp chem said NWChem was a good project.
 - The NWChem team lead didn't care that I didn't know how to program.
 - Argonne MCS told me I was useless to them if I didn't know C 😊
- The EMSL MSCF team to help me was the reason I was successful.
 - Karol Kowalski taught me coupled-cluster theory.
 - Bert de Jong taught me how to program in NWChem.
 - Dunyou Wang taught me how to use version control, manage big software.
 - Tim Carlson (sysadmin) taught me how to use Linux.
 - Many others... (I was a squeaky wheel)
 - I learned how to work as part of a team, build reusable software.

TL;DR

- I started working on NWChem in June 2006 and never stopped.
- My practicum project became my dissertation project.
 - Karol Kowalski mentored me remotely for three years.
 - CSGF funding gave me the freedom to do this and to travel to PNNL, ACS, etc.
- ...the basis for all of my postdoc offers and why I went to ALCF.
- NWChem has inspired essentially everything I have done since:
 - Trying to understand Global Arrays drove me into computer science.
 - Joined the MPI Forum and collaboration with the MPICH team.
 - Collaboration with Robert van de Gejin on dense linear algebra.
 - See the value in sustainable software, user support, etc.

Using CSGF skills in industry

- Teamwork, teamwork, teamwork.
 - Industry is a lot more like DOE than academia. *Rising tides lift all boats.*
 - Important projects have dozens if not hundreds of people.
 - While co-location is important to Intel, most teams are distributed.
- Interdisciplinary skill set
 - Application experts work on 2+ domains (e.g. MD and ML; Chem and Auto).
 - Ubiquitous use of mathematical modeling to understand performance, power.
 - Machine learning is applied mathematics...
- **Communication skills** are absolutely essential
 - Others are very smart, but are expert at something (very) different.
- *Software that works* is a primary currency

My day job

- Data Center Group (servers)
 - Enterprise and Government (i.e. not commercial cloud)
 - Exascale Co-Design Group
- Direct engagement with customers to drive development of:
 - Code modernization in their applications
 - Molecular dynamics, electronic structure (DFT), hydrodynamics (!)
 - Hardware features and system architecture for 3-7 years from now
- Continuous improvement of HPC software ecosystem
 - e.g. MPI, MKL, Fortran, OpenMP, and many OSS efforts
 - Collaboration with DOE on code modernization for KNL, OPA, ...
 - “Think like a user”



Jeff Hammond

jeffhammond

[Add a bio](#)

@Intel

Portland, Oregon

jeff.science@gmail.com

<http://jeffhammond.github.io/>

Organizations



Overview Repositories **96** Stars **322** Followers **141** Following **61**

Pinned repositories

[Customize your pinned repositories](#)

HPCInfo

Information about many aspects of high-performance computing. There is content in the Wiki.

C++ ★ 58 🍷 19

ParRes/Kernels

This is a set of simple programs that can be used to explore the features of a parallel platform.

C ★ 51 🍷 23

BigMPI

Implementation of MPI that supports large counts

M4 ★ 10 🍷 9

oshmpi

OpenSHMEM over MPI-3

C ★ 3 🍷 1

pmodels/armci-mpi

An implementation of ARMCI using MPI one-sided communication (RMA)

C


pmodels/casper

Process-based Asynchronous Progress Model for MPI Communication

C ★ 3 🍷 3

1,014 contributions in the last year

[Contribution settings](#)



[Learn how we count contributions.](#)

Less More

Intel-CSGF connections

- Intel is the largest industry employer of CSGF
 - Richard Mills and I talk regularly about PETSc, MPI, etc.
 - David Ozog recently joined Intel HPC programming models team
- IPCC program (funding for computational science)
 - NWChem teams at PNNL (Karol et al.) and LBNL (Bert et al.)
 - Devin Matthews at University of Texas – tensors and quantum chemistry
 - Ahmed Ismail at RWTH (now UWV) – LAMMPS molecular dynamics
 - David Keyes et al. at KAUST – scalable solvers
- Hal Finkel at Argonne is a key collaborator on HW-SW co-design

...I'm not the only one

 **Jesse Lopez** 10:19 am (1 day ago)   

to me 



My current industry employer (mostly a government contractor) absolutely loves the perspective and experience with HPC I brought to the company. CSGF is to thank for that.

Thanks again,
-jesse

Summary

- The CSGF program is the single most important thing that has happened to me in my *professional* life.
- The mentors, mentees, and collaborators I met in CSGF continue to enrich my work in computational chemistry and linear algebra.
- **Intel does not hire computational chemists with mediocre Matlab skills.**

Questions