NSF PI Meeting: The Science of the Cloud, Mar 17-18, 2011
Waterview Conference Center, Arlington, VA

Cloud Computing Clusters for Scientific Research*

J. J. Rehr, F. D. Vila, and K. Jorissen

Department of Physics, University of Washington
Seattle, WA 98195-1560

*Supported by NSF SI2-SSE Grant OCI-1048052

Challenge of NSF Grant

- Complete scientific computing cloud environment: for materials science, robust, easy to use

Goals of CC for Scientific Computing*

- On-demand access without the need to purchase, maintain, or understand HPCs
- Optimized/pre-installed HPC scientific codes
- Low cost HPC access to a wide class of scientists

*J. J. Rehr, F. Vila, J. P. Gardner, L. Svec and M. Prange
CiSE May/June 2010

The Target Scientific CC User

- Not usually HPC “gurus”:
  - Little experience with parallel computing, code compilation & optimization, etc.
  - Limited access to HPC resources
- “Would rather be doing science...”

Scientific Cloud Computing Programming Philosophy

- Preserve current HPC paradigm (MPI) on the cloud
- Simplify access to HPC for a broad user base

Development Strategy

1. Develop AMIs (Amazon Machine Images) for HPC scientific applications in materials science
2. Develop shell-scripts that make the EC2 look & run like a local HPC cluster: ”virtual supercomputer on a laptop”
3. Test serial, parallel & network performance
4. Develop GUI to control/run codes & I/O

Cloud→Cluster with UW Cluster tools*

User interacts with control workstation

Control Workstation

ec2-clust-*

1. Start cluster
2. Configure nodes

MPI Master

startd

MPI Slave

startd

MPI Slave

startd

EC2 Compute Instances

* http://www.phys.washington.edu/feff/scc

UW EC2 Cluster Tools

UW EC2 tools on local control machine

Name	Function	Analog
`ec2-clust-launch N`	Launches cluster with N instances	`boot`
`ec2-clust-connect`	Connect to a cluster	`ssh`
`ec2-clust-put`	Transfer data to EC2 cluster	`scp`
`ec2-clust-get`	Transfer data from EC2 cluster	`scp`
`ec2-clust-list`	List running clusters	
`ec2-clust-terminate`	Terminate a running cluster	`shutdown`

The tools hide a lot of the “ugliness”:

`ec2-clust-connect`

```
ssh -i/home/fer/.ec2_clust/.ec2_clust_info.7729.r-de70cdb7/key_
pair_fdv.pem feffuse@ec2-72-44-53-27.compute-1.amazonaws.com
```

Current EC2 AMIs

- New 32-bit and 64-bit AMIs
- Amazon's HPC "cluster instance"
- EBS backed instances (faster disk access/boot time)

- Several Condensed Matter/Materials Science Codes:
  - X-ray Spectroscopy: FEFF9, IFEFFIT
  - Electronic Structure: WIEN2k, ABINIT, Quantum Espresso
  - Excited States: AI2NBSE, OCEAN, Exc!ting, RT-SIESTA

UW CC Tools: 2nd generation (v2)

- Older v1 tools :
  - Emphasized simplicity & maximum security
  - SLOW: 12 min for n_{inst}=16, linear with cluster size

- v2 Solutions:
  - Asynchronous tools: minimal setup, max. security
  - Self-Assembling clusters: fast, but sacrifices security
  - Current: Hybrid approach; non-sensitive tasks on cloud keeps credentials safe

UW v2 Tools: Boot and Setup Times

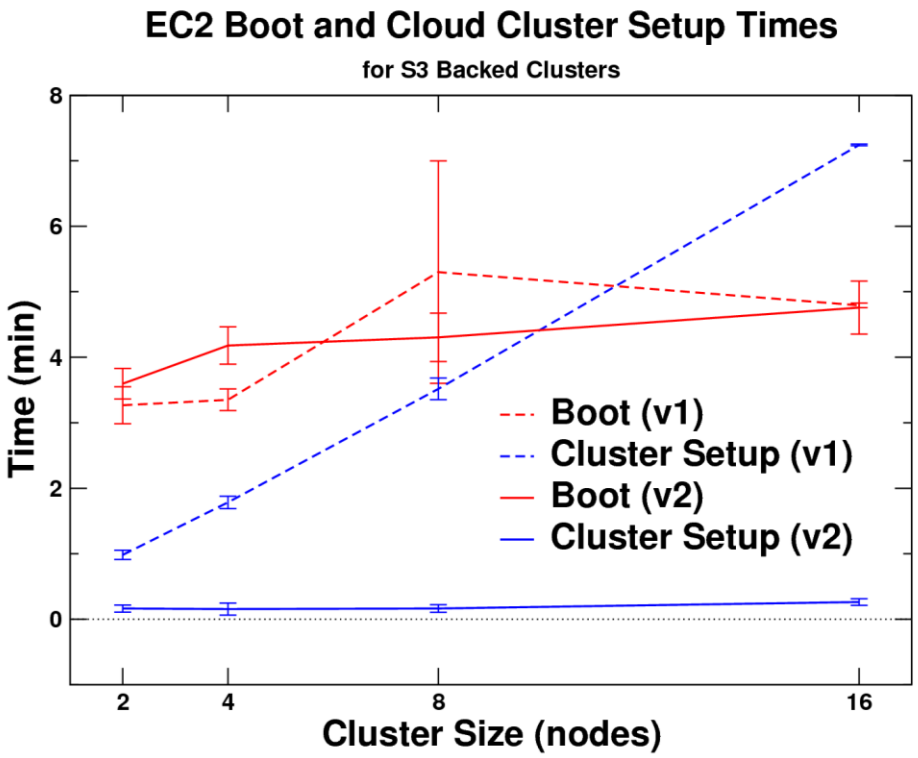

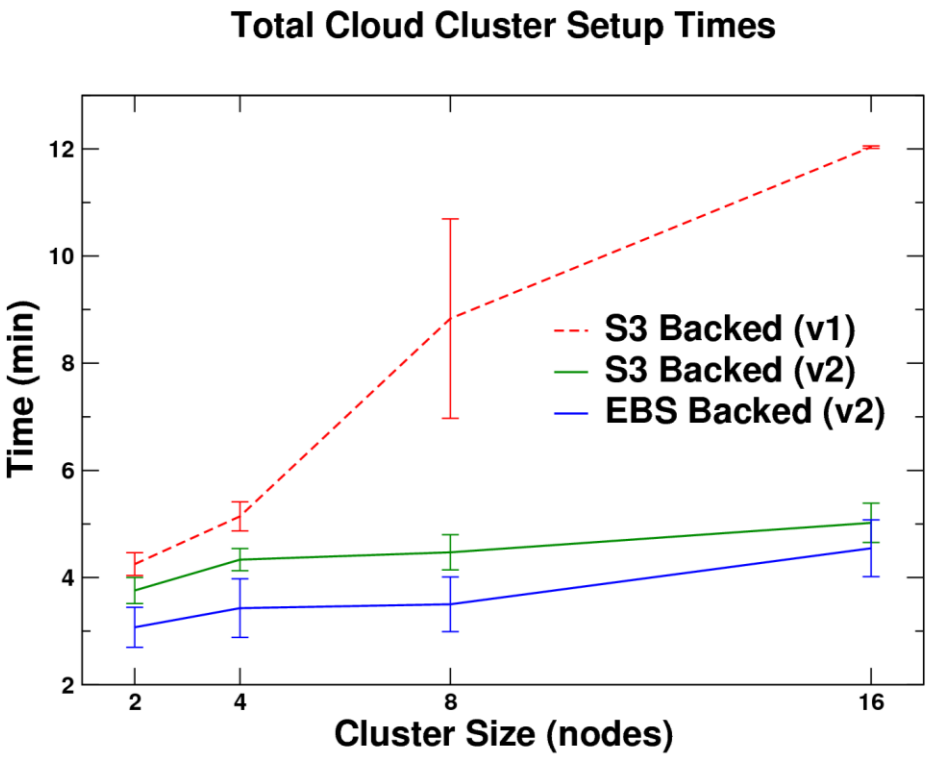

- Setup time negligible & independent of cluster size!
- EBS backed clusters faster than S3

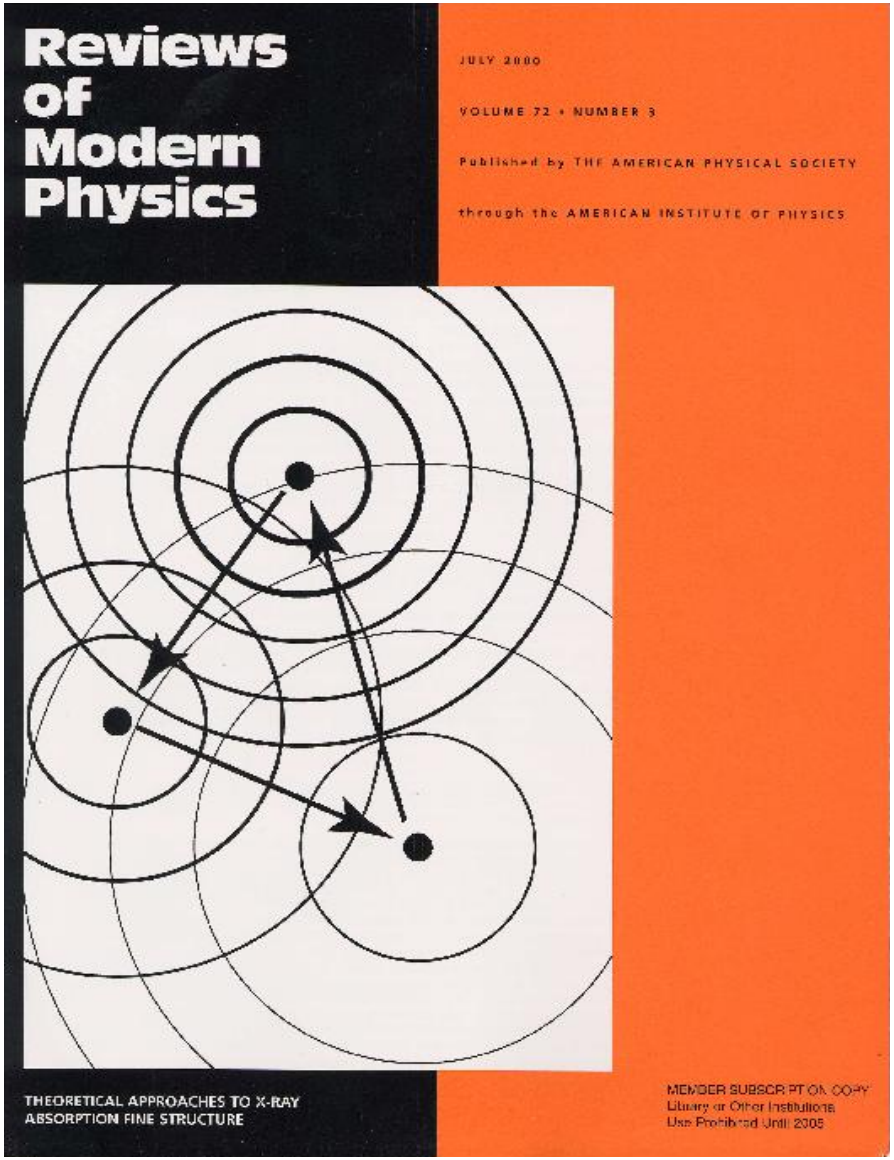

Sample Scientific Application

FEFF:

Real-space Green's function code for electronic structure, x-ray spectra, EELS, …

User Base:

1000s of users in physics, chemistry, materials science, biophysics…

J. J. Rehr & R.C. Albers

Rev. Mod. Phys. **72**, 621 (2000)

http://leonardo.phys.washington.edu/feff/

FEFF Parallel Performance

Naturally parallel: Each CPU calculates a few points in the energy grid

Loosely coupled: Little communication between processes

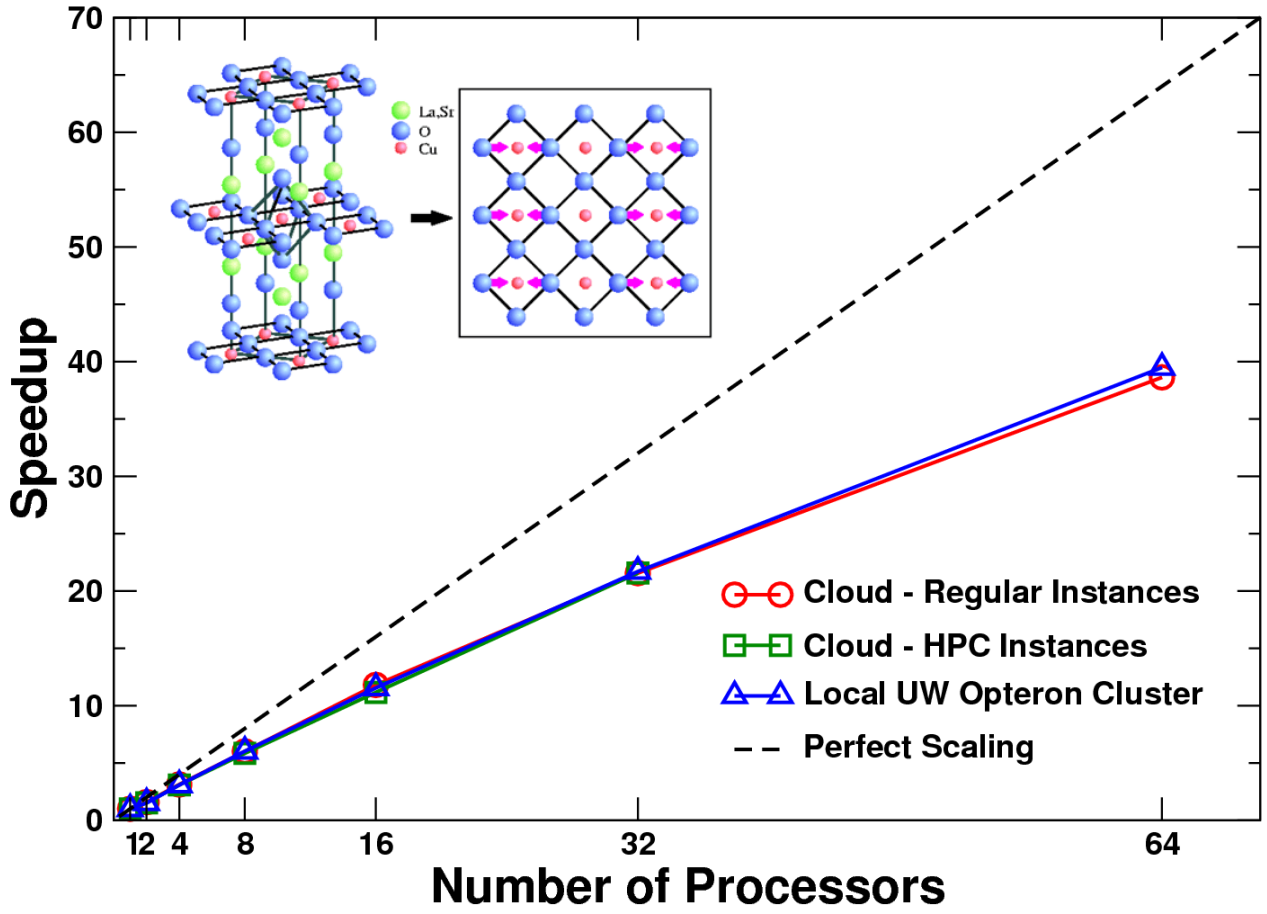

- FEFF9 shows identical speedup on local cluster and EC2 cluster

Sample Electronic Structure Application WIEN2k Parallel Performance

DFT KS equations on 128 k-point grid

Time to solve one k-point in one core

EC2 small	312s
EC2 medium	116s
EC2 HPC	84s
UW Opteron	189s

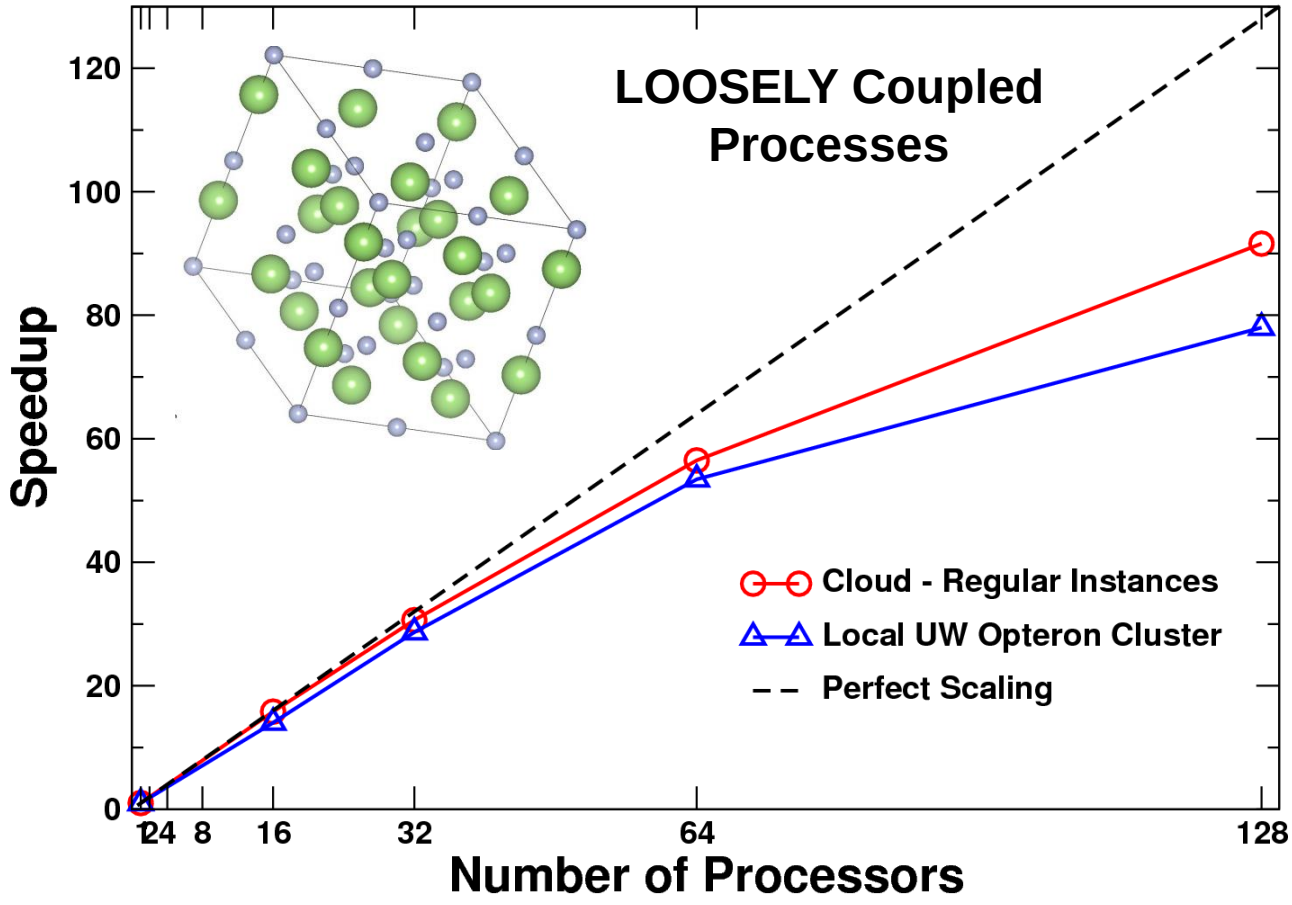

- Good scaling in regular instances

WIEN2k Parallel Performance

KS for large system at 1 k-point

VERY DEMANDING of network performance

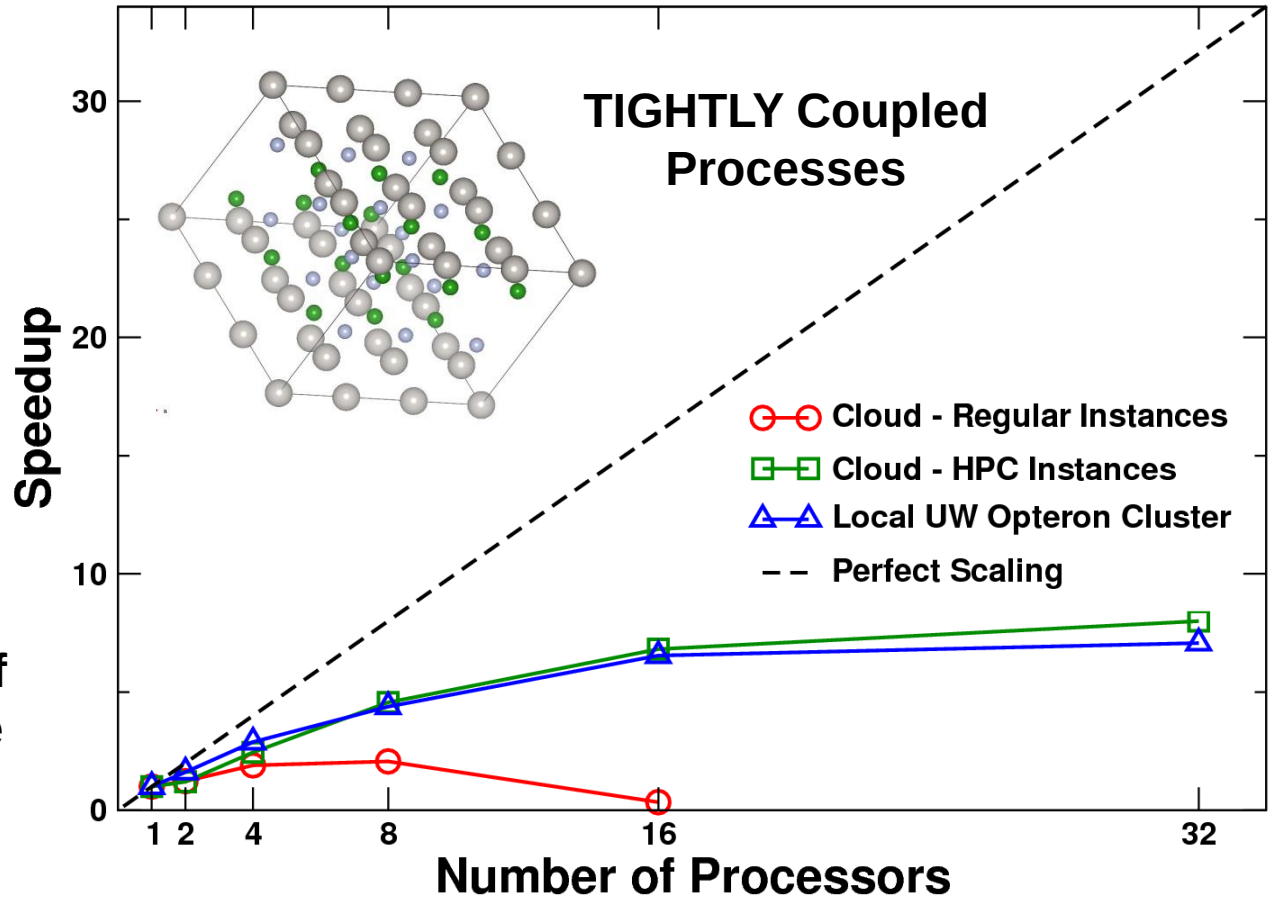

- Regular instances cannot deliver network-intensive performance
- HPC cluster instances deliver the same speedup as local Infiniband cluster

FEFF GUI: XAS on the Cloud

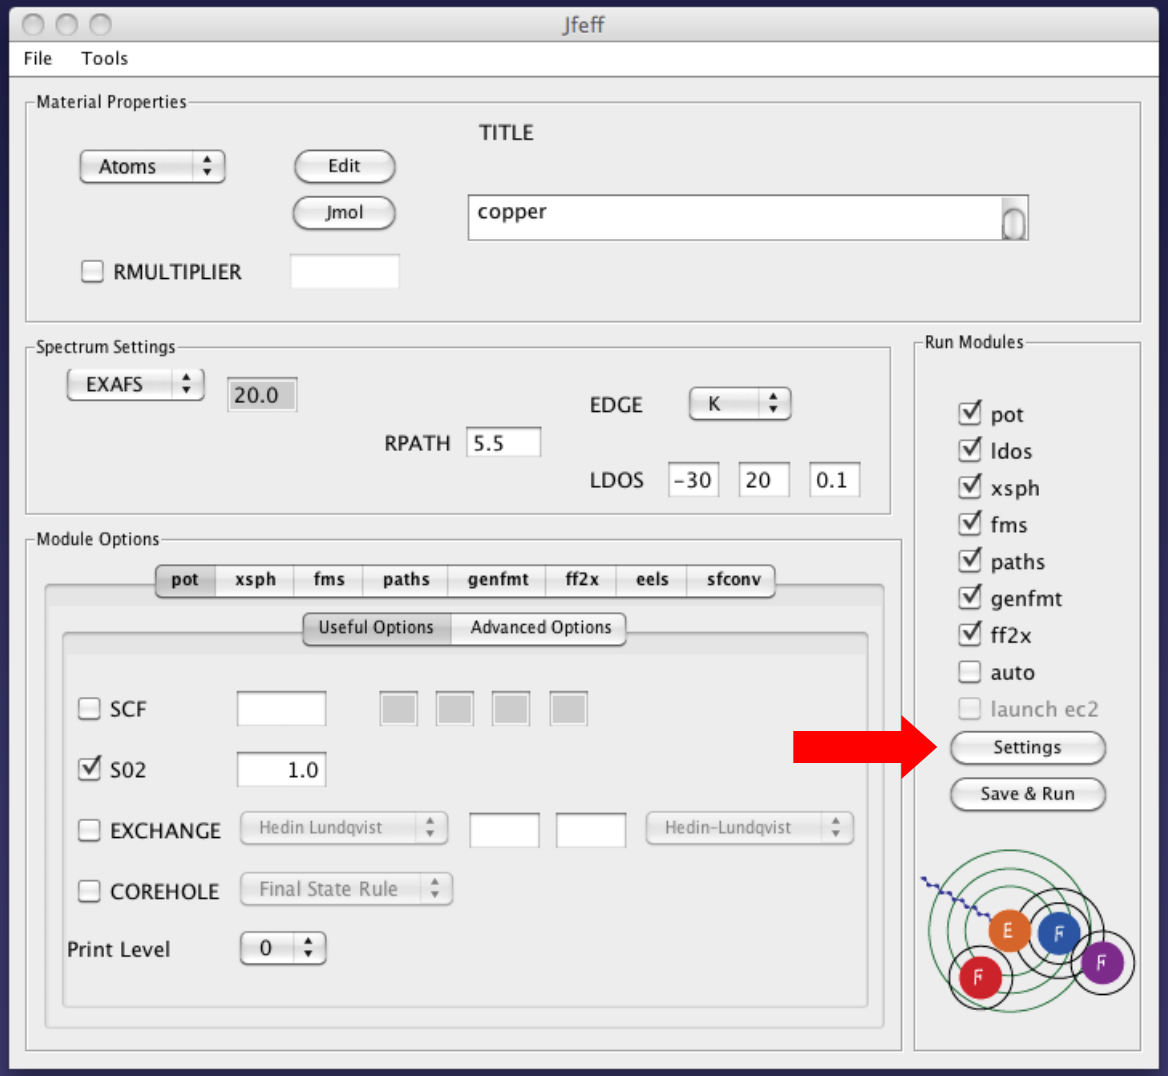

FEFF GUI: XAS on the Cloud

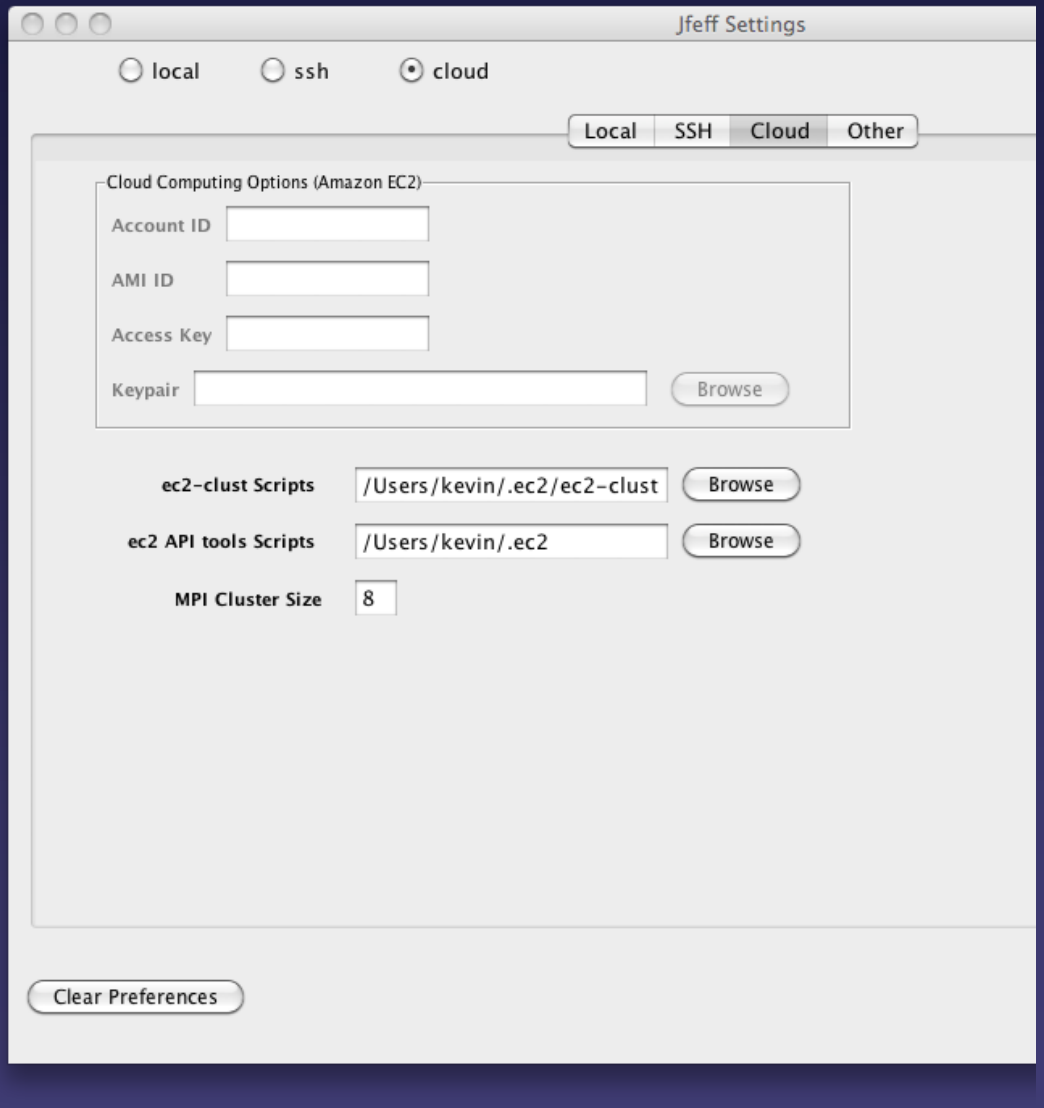

FEFF GUI: XAS on the Cloud

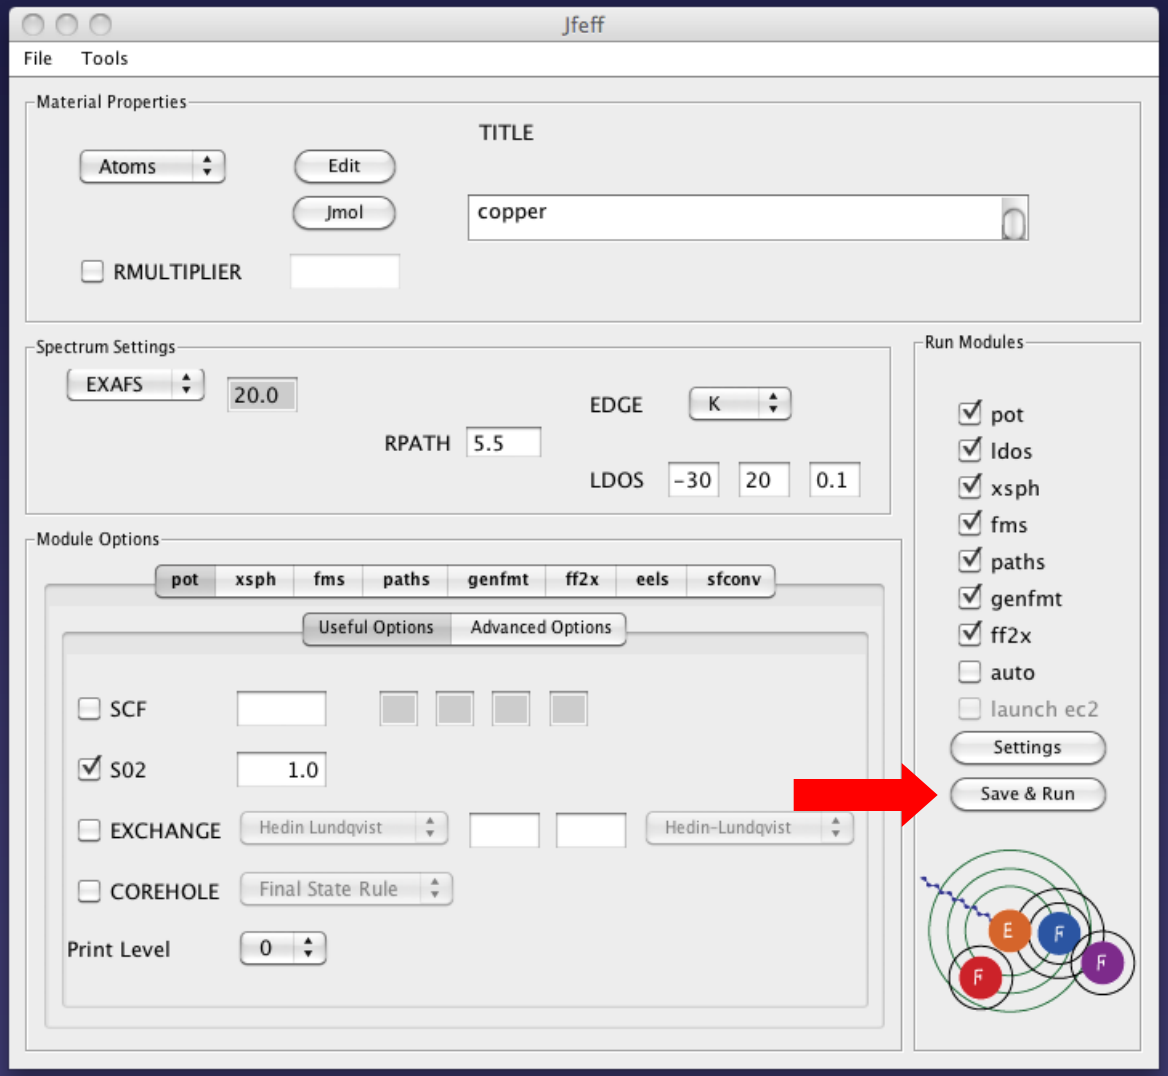

FEFF GUI: XAS on the Cloud

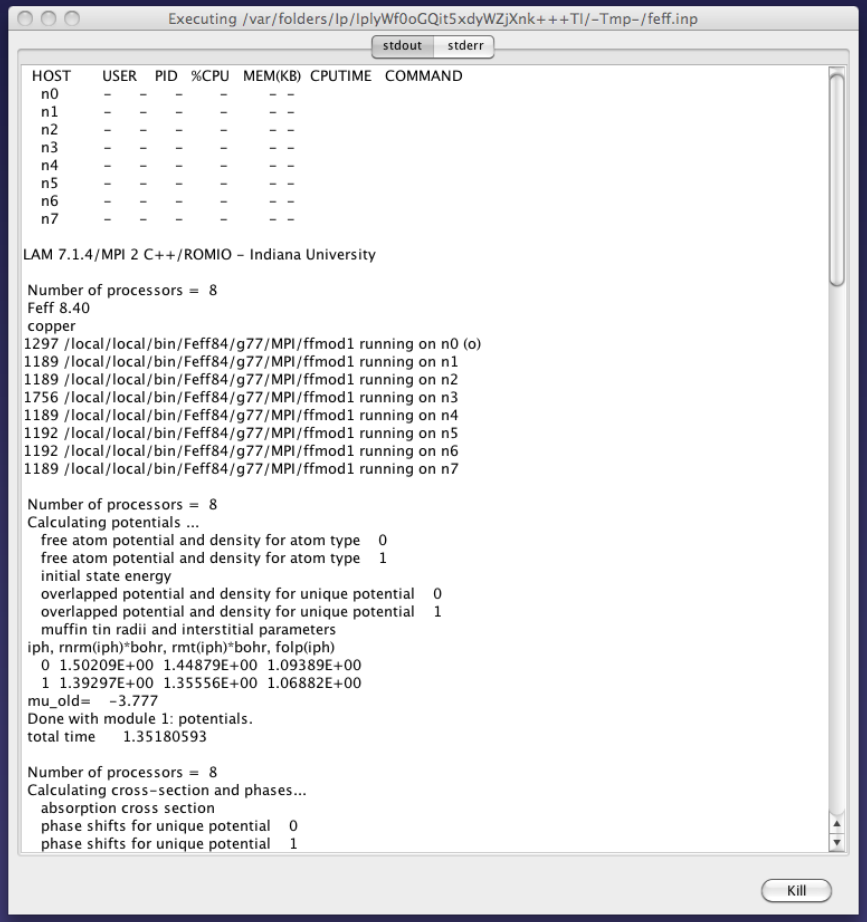

FEFF GUI: XAS on the Cloud

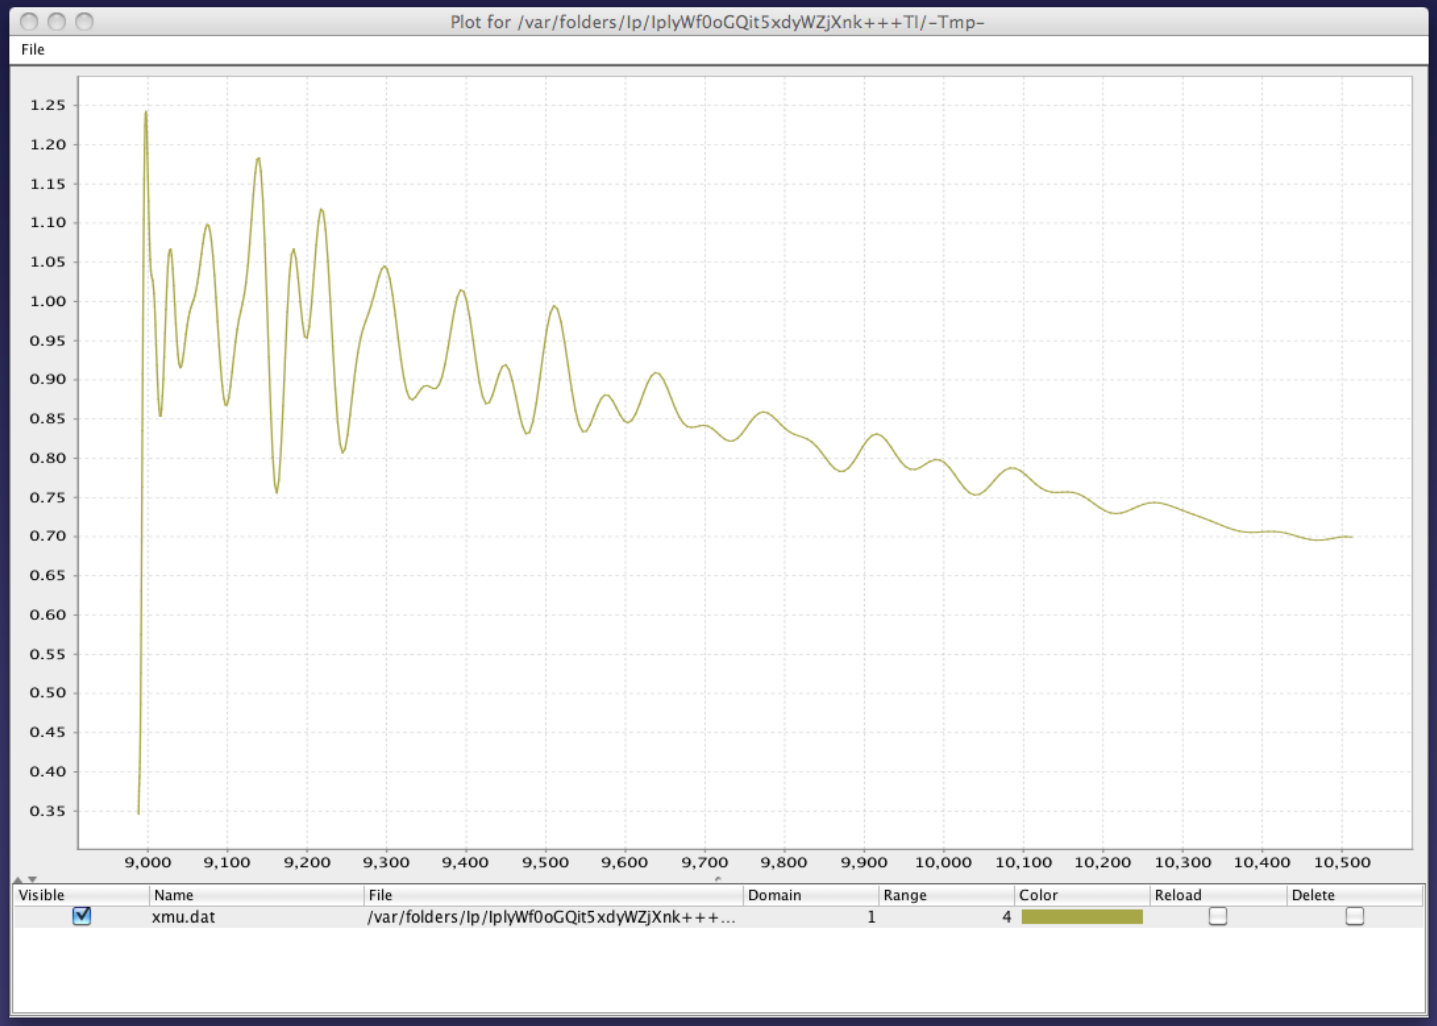

Conclusions

- Goals partially achieved: robust SCC environment developed & benchmarked on FEFF, WIEN2k
- SCC-ready FEFF App with JFEFF GUI (OSX, Linux)
- Improved SCC AMIs and UW Cluster tools (v2):

  Regular EC2 AMIs OK for light network use

  HPC AMIs NEEDED for MPI+ScaLAPACK

Future Project Goals

- Develop & add WinPC support
- Develop SCC Java GUI for CC management, runs, I/O
- Optimize & benchmark other advanced scientific codes
- Distribute to Scientific Community: http://www.phys.washington.edu/feff/scc

Acknowledgments

UW: L. Svec, J. Loudermilk, D. Bitseff. M. Prange, J. Gardner, R. Coffey, E. Lazowska

Amazon: D. Singh, J. Barr, T. Laxdal, P. Sirota, P. Sridharan, R. Valdez, and W. Vogels

NSF: C. Bouldin

UW-SCC is supported by NSF grant OCI-1048052

FEFF is supported by DOE-BES grant DE-FG03-97ER45623

... and thank you