Penn State, August 2013

Cloud-WIEN2k

A Scientific Cloud Computing Platform for Condensed Matter Physics

K. Jorissen

University of Washington, Seattle, U.S.A.

Supported by NSF grant OCI-1048052

www.feffproject.org

Materials Science

Materials Science research:

Theoretical models, evaluated on a computer, are usually needed for interpretation and quantification of measurements.

But HPC is often not readily available.

sample

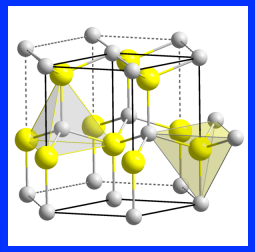

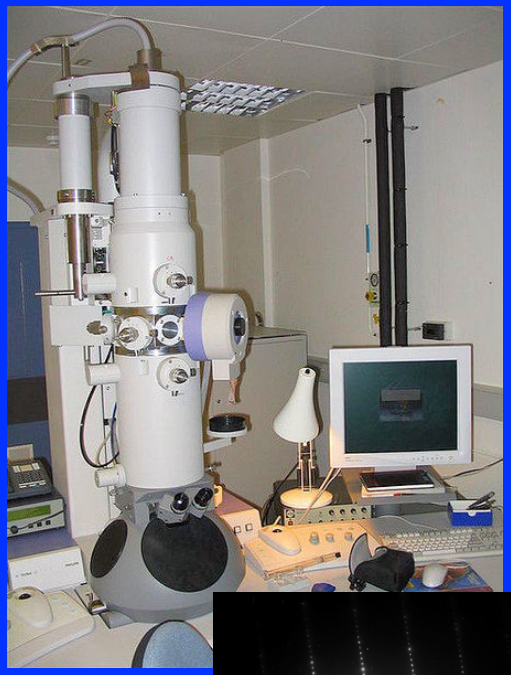

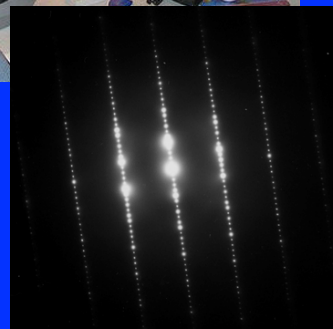

$H\psi=E\psi$

$E=mc^2$

theoretical model

measurement

interpretation

Anecdote *(High-Performance Computing is everywhere)*

Computational linguistics:

"We automatically identify semantically related words in the 400 million word Dutch Twente corpus to

→ Statistically find contextual associations and quantify association strength

→ Identify syntactical relations between words

→ Relevant to automatic translation software

Multivariate analysis with dozens of variables – large computational needs."

--- an "English Lit major"

https://perswww.kuleuven.be/~u0042527/lingpub.htm

How do we bring the best theory and simulations to the scientists who need it?

(often applied scientists – not computational specialists)

SOLUTION:

Scientific Cloud Computing

Are state-of-the-art calculations "work for specialists"?

FEFF-old (simple Einstein model for phonons)

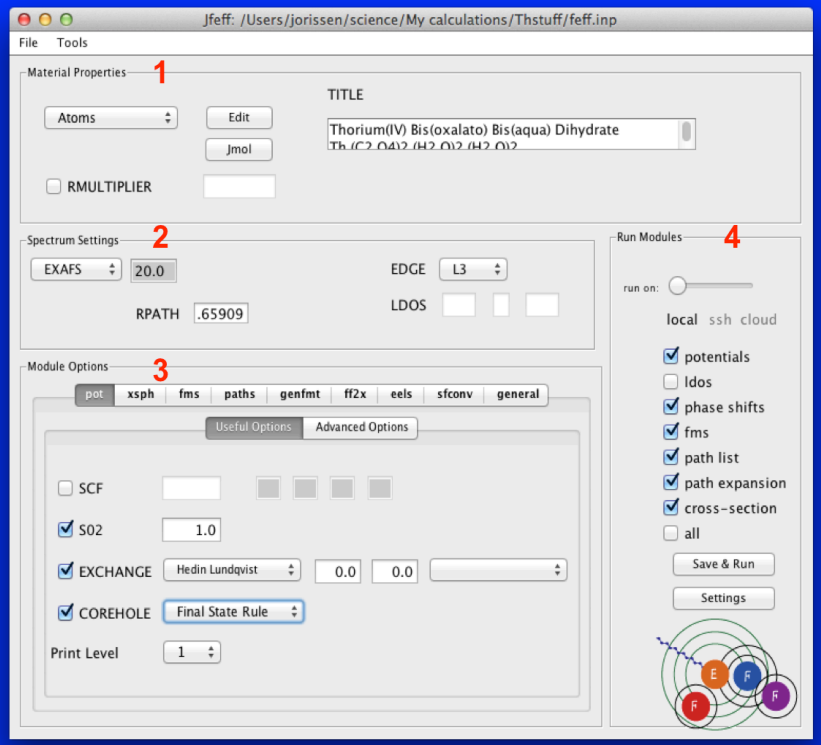

GUI
Easy install
Runs on laptop
Load file & Click "Run"

~ 1 day to learn

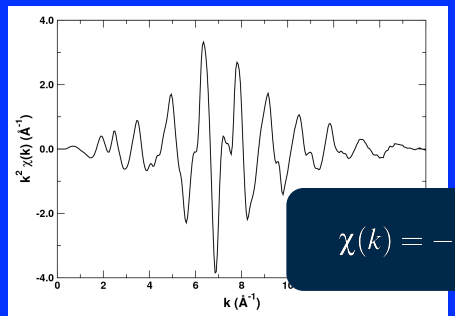

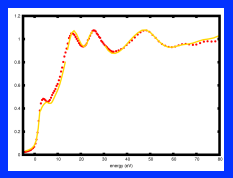

FEFF

$$\chi(k) = -k\sum_i S_0^2 N_i \frac{f_i^{eff}(k)}{R_i^2} e^{-R_i/\lambda} \sin\left(2kR_i + 2\phi(k)\right) e^{-\sigma_i^2 k^2/2}$$

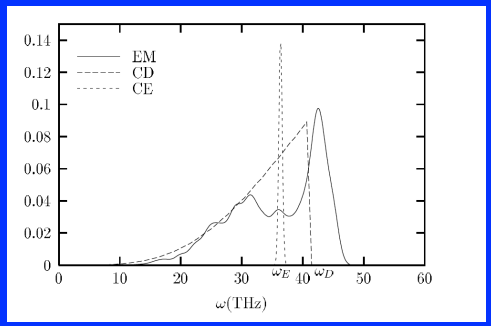

Are state-of-the-art calculations "work for specialists"?

FEFF-gold (accurate ab initio model for phonons)

Dyna

DFT requires cluster

Deby

Invented / published 2006-2009
Clearly an improvement
Nobody uses it

X-ray Absorption -- FEFF

~ 0.x grad students to learn

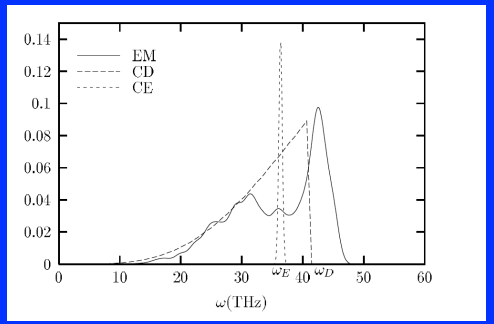

Are state-of-the-art calculations "work for specialists"?

- Hardware barrier: advanced codes need clusters
  - Buy a cluster? IT support?
  - Supercomputing center?
  - Collaborate with specialists?
- Software barrier: running codes is difficult
  - Installation of || software tricky
  - lacking user-friendliness
  - multi-code workflows difficult

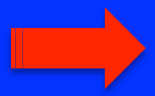

t >> 1 before improved theory reaches applied research

SC2IT
Scientific Cloud Computing

Scientific Cloud Computing

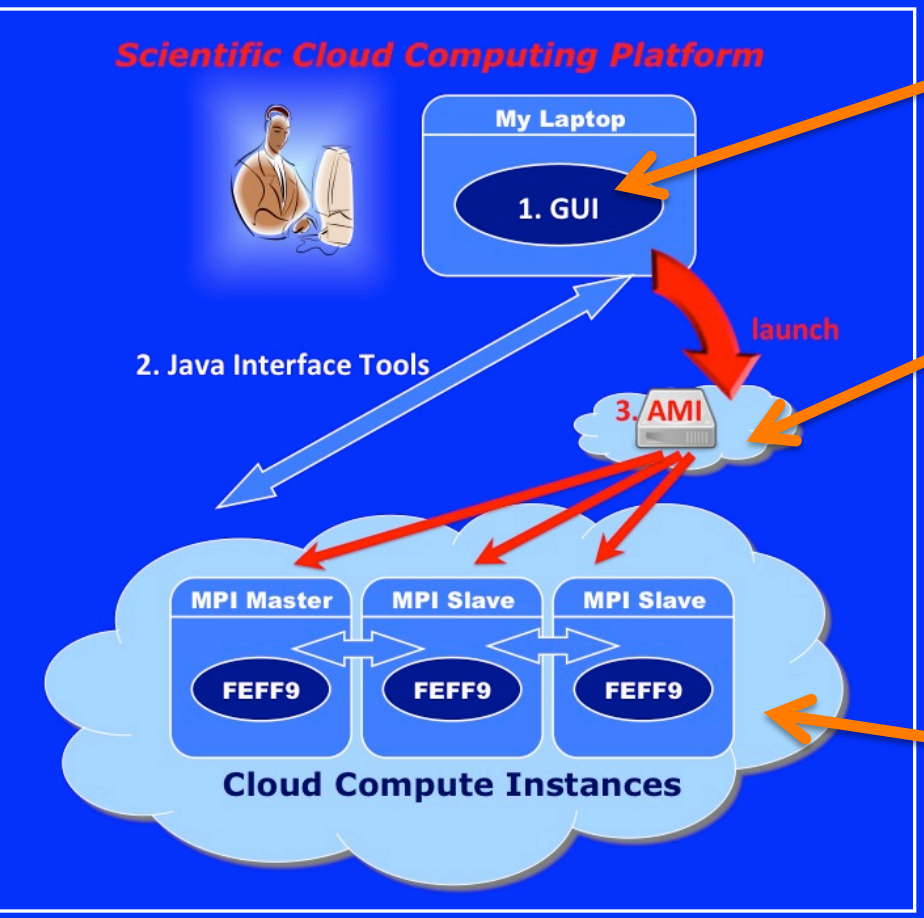

Interface simplifies workflow (hides cloud -- app)

Developer makes virtual "XAS" compute node with preinstalled WIEN2k

User requests 5 node Cloud Cluster for 3 hours when needed ($20)

amazon web services™

SC2IT
SCIENTIFIC CLOUD COMPUTING

SCC Virtual Machine Image

- Contains utilities for parallel scientific computing:
  MPI, compilers, libraries, NFS, …
- Becomes compute node in SCC Cloud Cluster

- Developer-optimized Scientific codes for your research
  - WIEN2k for electronic structure calculations
  - latest version
  - optimized for performance
  - MPI parallellization for large calculations

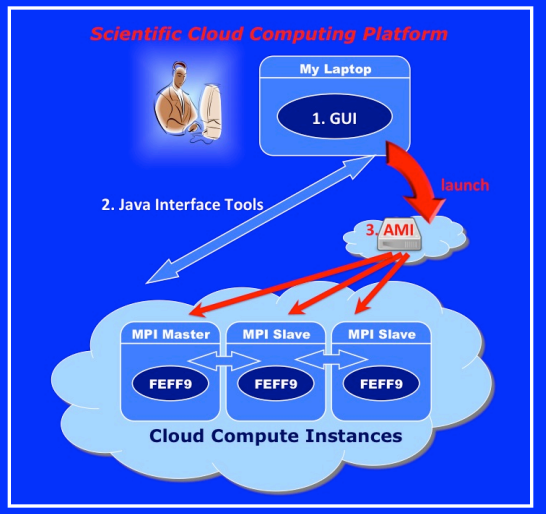

"My new research group was looking for a way to implement *MEEP-mpi* (MIT Electromagnetic Equation Propagation) to simulate EM fields in nanoscale optical devices for cavity QED experiments. ***We believe that Amazon EC2 is an economical and time saving solution for our finite difference time domain (FDTD) simulations.*** My group's research iterates between fabrication and simulation thus it is advantageous to buy computing power only when needed. Moreover ***it is a relief not to have to maintain our own small cluster within our group.***"

Kai-Mei Fu, University of Washington (USA)

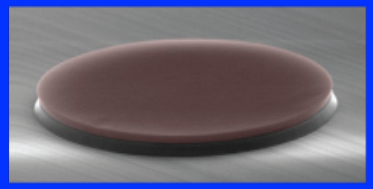

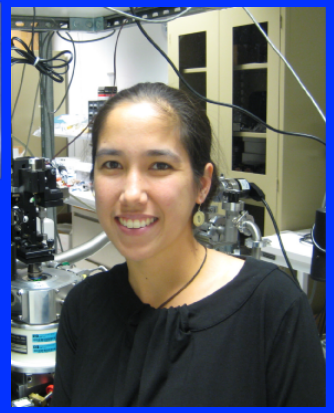

SC2IT
Scientific Cloud Computing

SCC Linux interface

For developers of GUIs

Java interface library (jar)

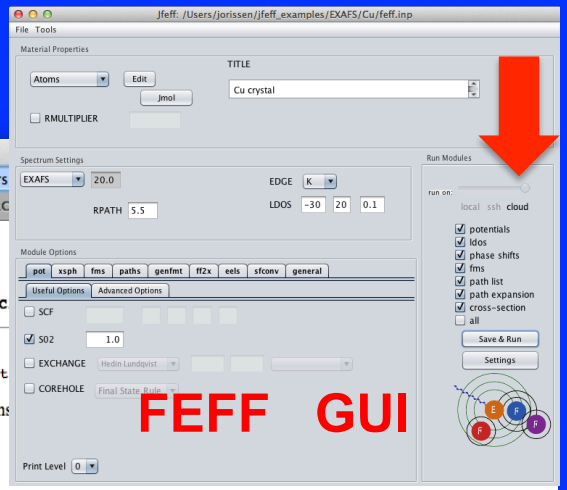

Class Cluster

```
java.lang.Object
  └ edu.washington.scc.C
```

```
public class Cluster
extends java.lang.Object
```

The primary class that contains

Author:
William Johnson

Constructor Summary

`Cluster(java.lang.String name, ClustAuth auth)`	Constructs a new Cluster object with the given name and using the given credentials.

Method Summary

ClusterResult	`connect(boolean toRoot)` Opens a Swing window which contains a terminal on the remote machine.
ClusterResult	`executeCommand(java.lang.String cmd)`
ClusterResult	`executeCommandRoot(java.lang.String cmd)` Executes a single bash command via ssh on the remote machine.
ClusterResult	`get(java.lang.String remoteName, java.lang.String localName)` Downloads a file from head node of the cluster to the local machine.
java.util.Properties	`getConfig()` Returns a Properties object which contains all pertinent configuration inform
java.lang.String	`getName()`
ClusterStatus	`getStatus()`

SCC Java interface

For savvy users and developers

Collection of shell scripts

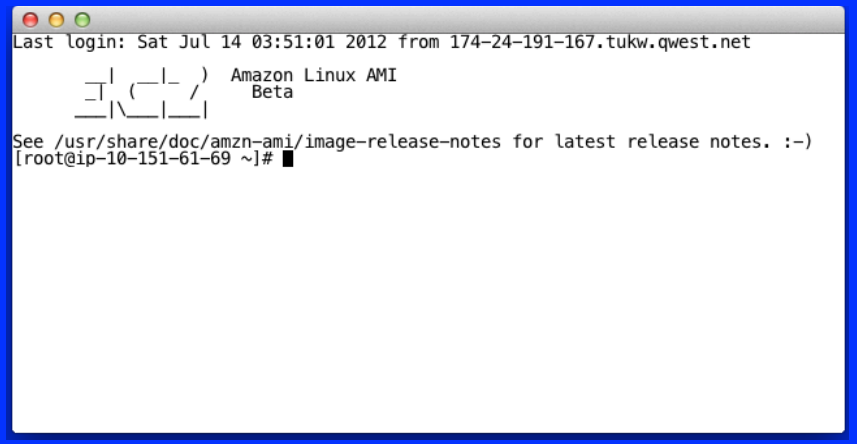

Tools in the SC2IT interface for bash

Name of tool	*function*	*analog*
sc2launch N	Launch cluster with N instances	boot
sc2connect	Connect to a cluster	ssh
sc2put	Transfer data to a cluster	scp
sc2get	Transfer data from a cluster	scp
sc2list	List running clusters	top
sc2terminate	Terminate a running cluster	shutdown
sc2desc	Describe all running clusters	ps
sc2run	Launch a job on a cluster	.
sc2connectr	Connect to a cluster as root	ssh
sc2load	Monitor load in a cluster	loadavg

WIEN2k-cloud

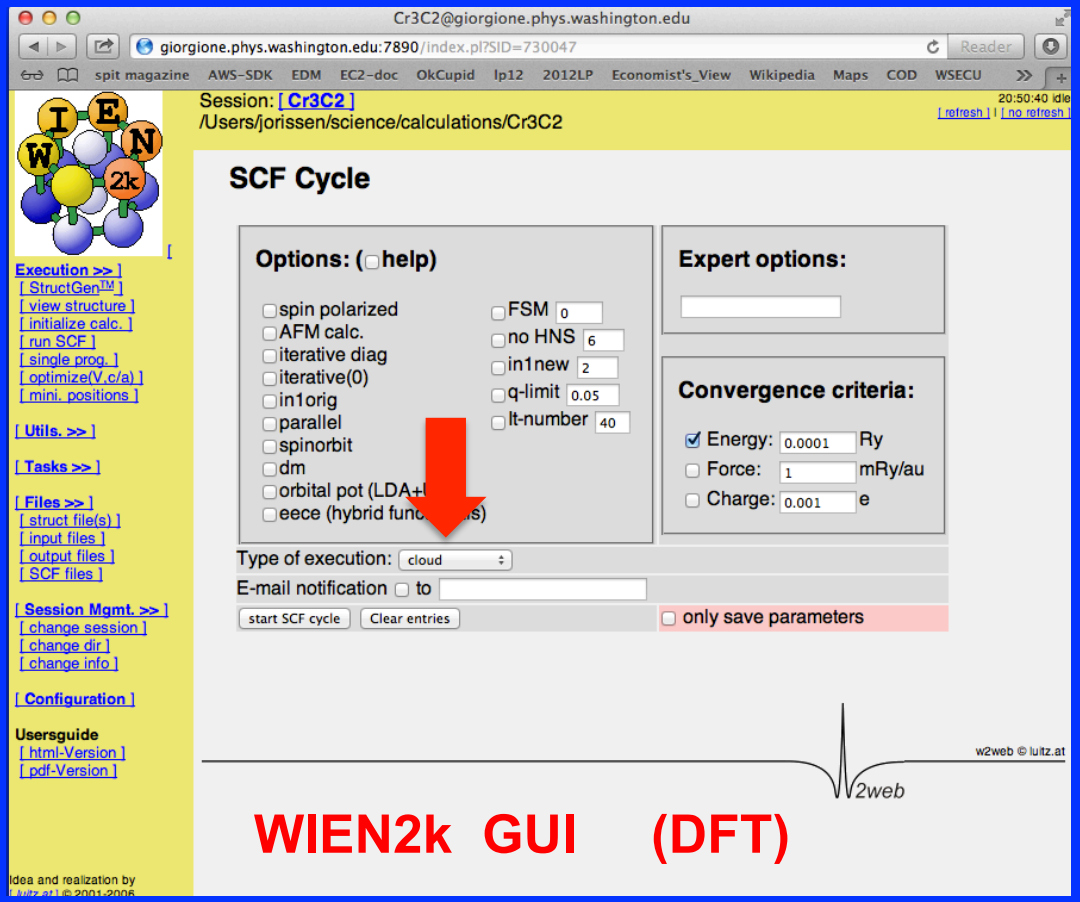

- Starts || cluster in EC2 cloud
- Uploads initialized calculation
- Runs || calculation in EC2 cloud
- Downloads results to laptop
- Deletes EC2 cluster

Other workflows / data flows can be added.

Requires:

- create EC2 account
- install SCC program

Performance

LOOSELY Coupled Processes

DFT KS equations on 128 k-point grid

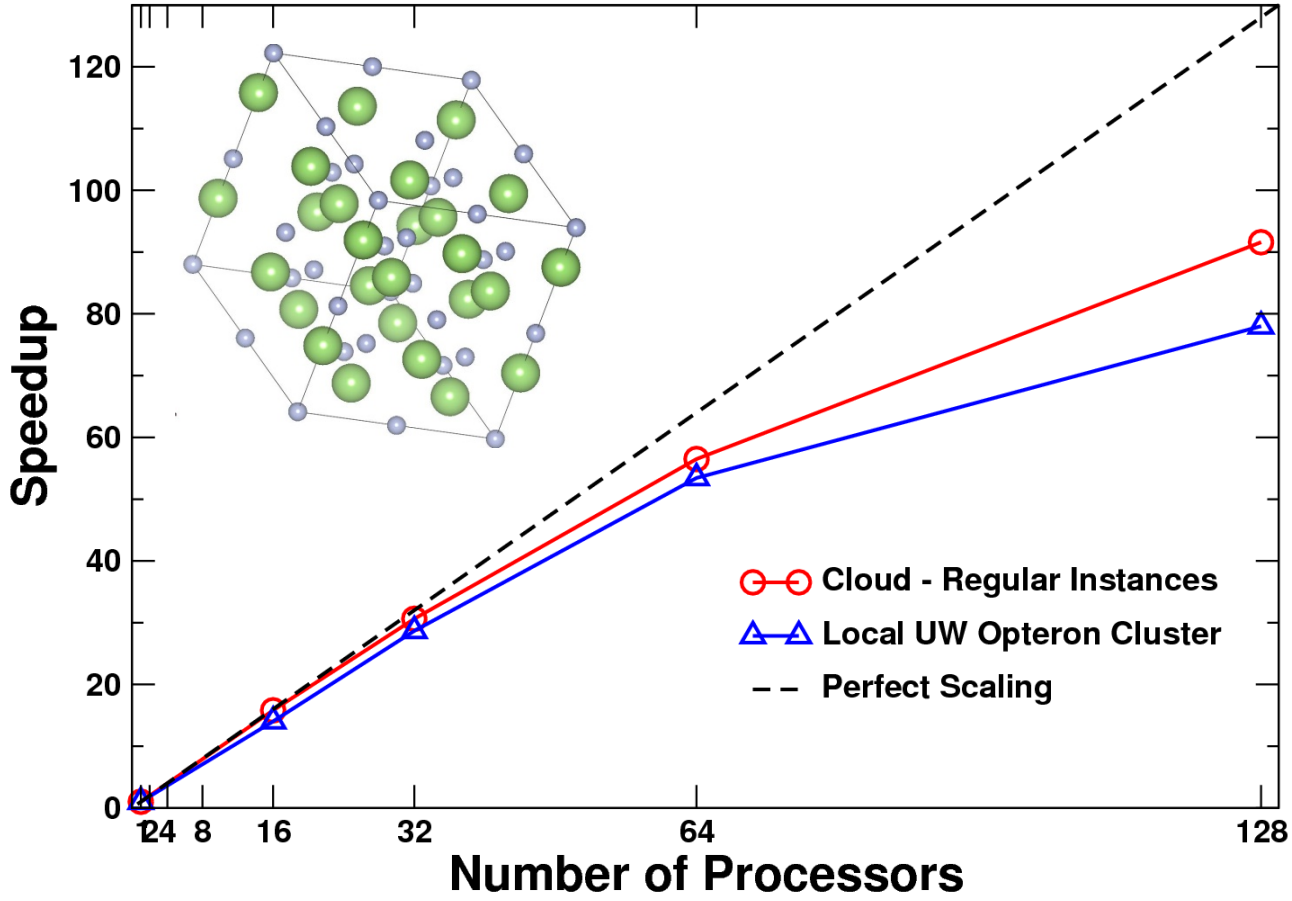

Good scaling

Performance

TIGHTLY Coupled Processes

KS for large system at 1 k-point

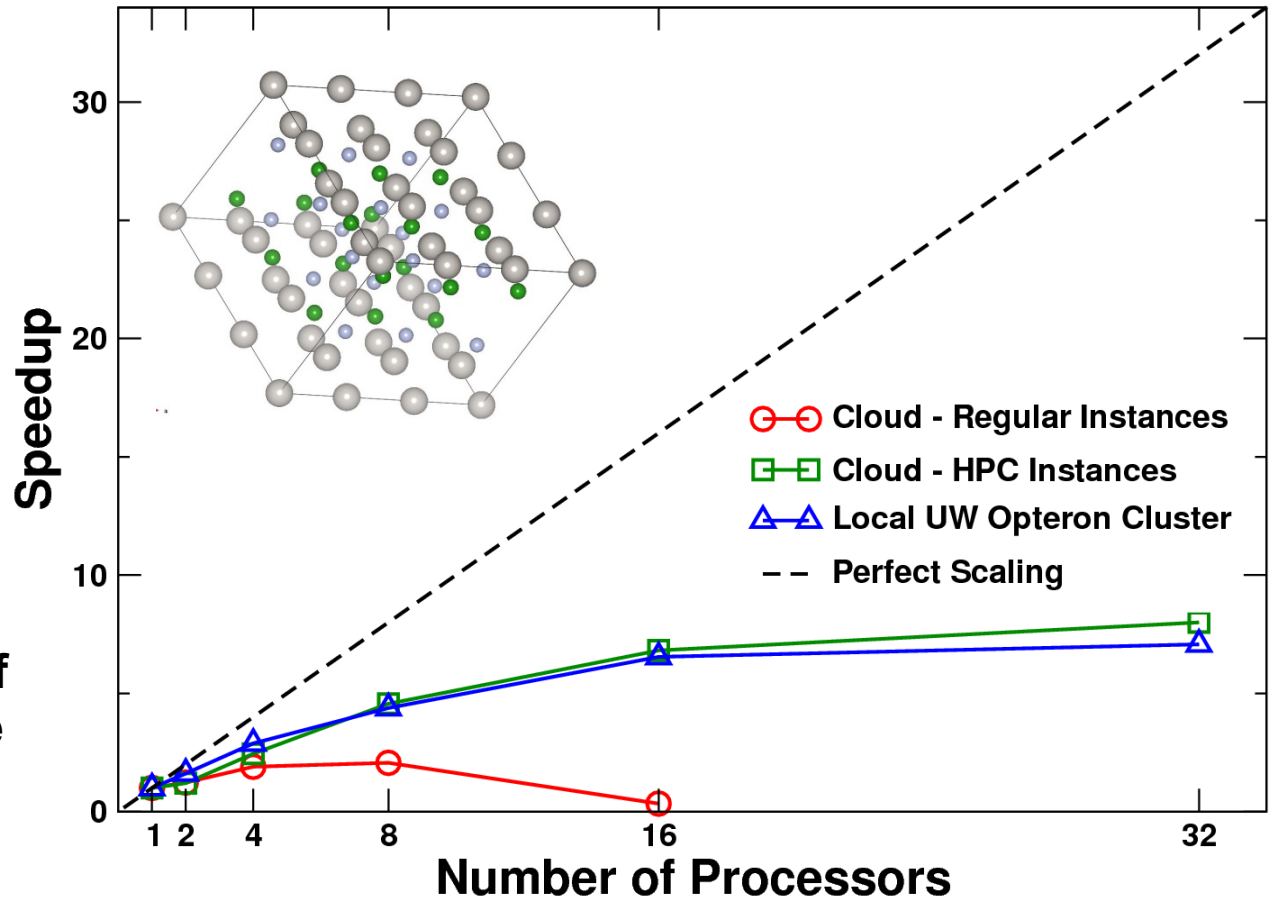

VERY DEMANDING of network performance

HPC cluster instances deliver good speedup

5. WIEN2k Performance Benchmarks

TIGHTLY Coupled Processes

KS for large system at 1 k-point

H size 56,000 (25GB)
Runtime (16x8 processors) :
Local (Infiniband) 3h:48
Cloud (10Gbps) 1h:30 ($40)

VERY DEMANDING of network performance

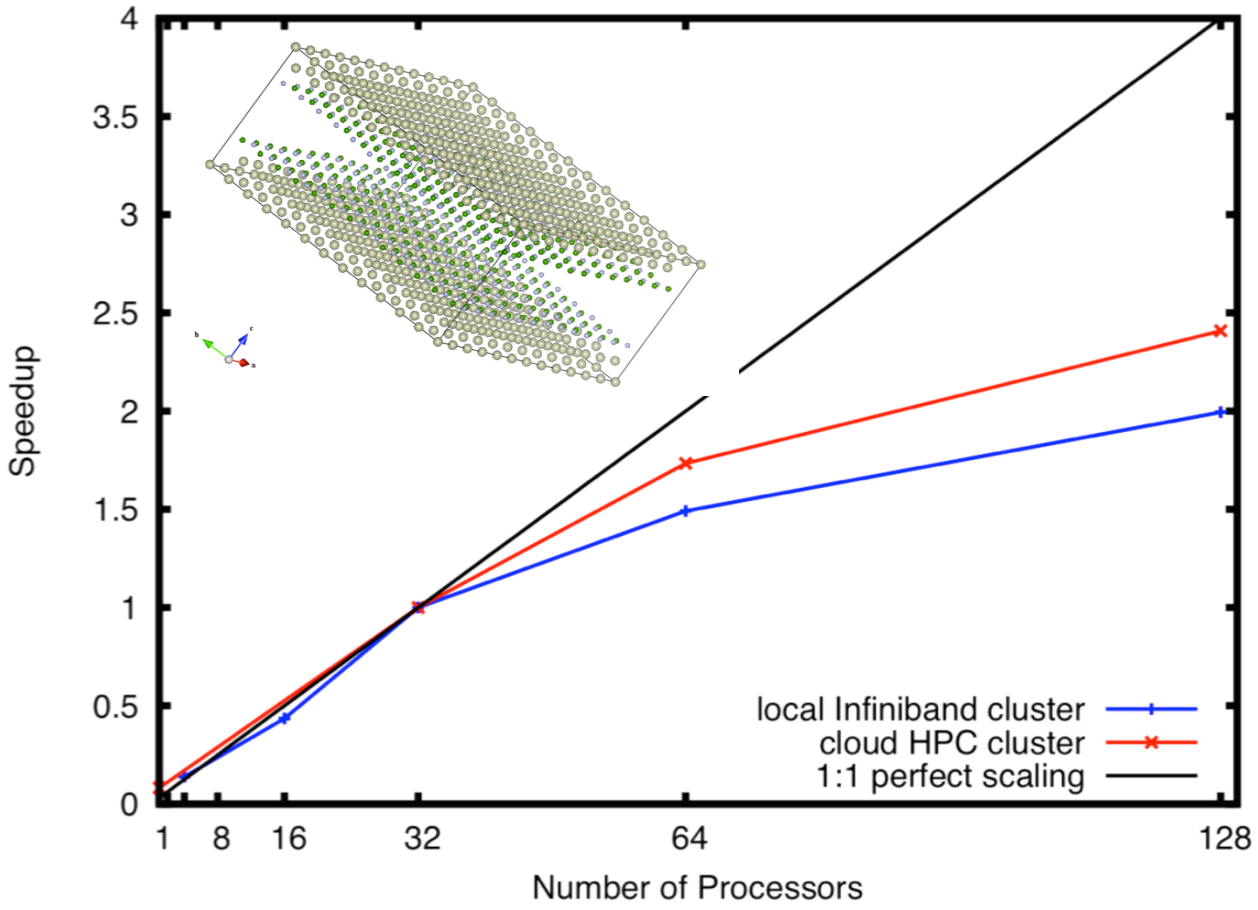

- 1200 atom unit cell; SCALAPACK+MPI diagonalization, matrix size 50k-100k
- HPC cluster instances deliver similar speedup as local Infiniband cluster

"Scientific Cloud Computing can bring novel theory & HPC modeling to more researchers."

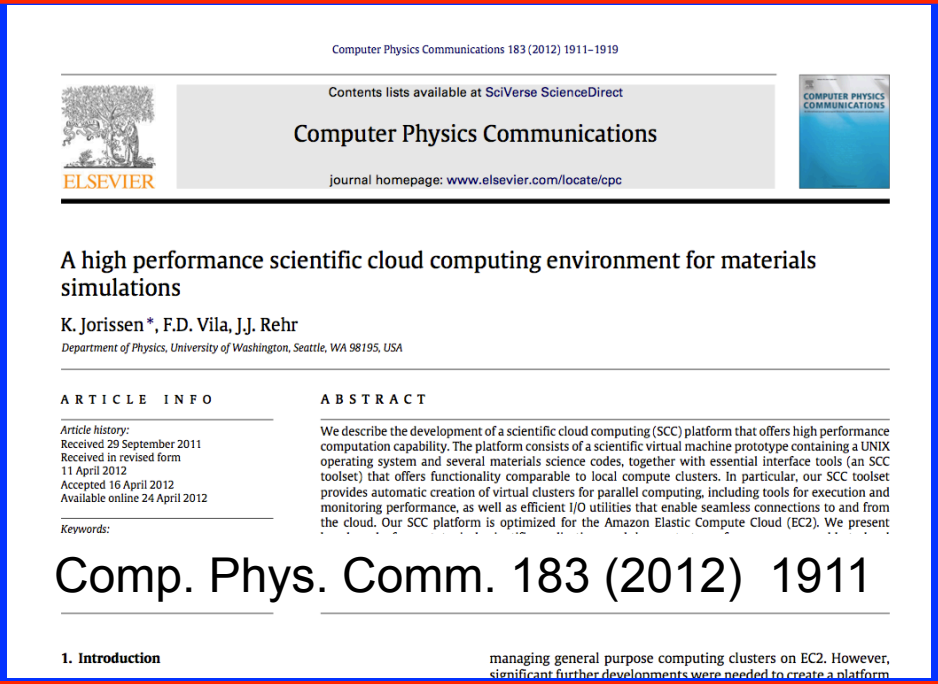

We acknowledge: ◆ **FEFF: S. Story T. Ahmed B. Mattern M. Prange J. Vinson ◆ UW: R. Coffey E. Lazowska J. Loudermilk ◆ Amazon: D. Singh ◆ NSF: C. Bouldin ◆ supported by NSF OCI-1048052**

Backup stuff

4. Cloud-Computing on the Amazon EC2 cloud

My Laptop

FEFF interface

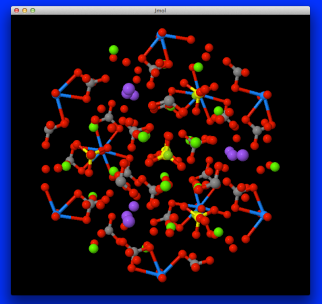

1. Create cluster
2. Calculations
3. Stop cluster

MPI Master

FEFF9

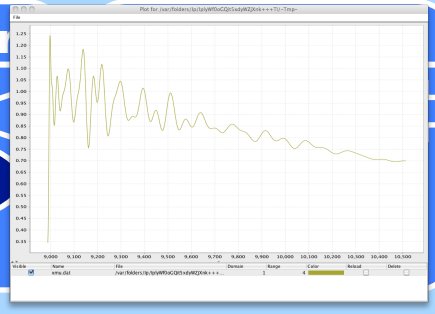

MPI Slave

FEFF9

Cloud Compute Instances

* K. Jorissen et al., Comp. Phys. Comm. 183 (2012) 1911

Developer's view:

ExecuteCloudContext.java:

```
import edu.washington.scc.*;

// Launch the new cluster with "cs" specifications:
ClusterResult rl = clust.launch(cs);

// Initialize the FEFF calculation on the cloud cluster:
// Copy feff.inp:
ClusterResult rp = clust.put(LocalWorkingDir+"/feff.inp", CloudWorkingDir+"/feff.inp");

// Run the FEFF9-MPI calculation:
ClusterResult rf9 = clust.executeCommand(Feff9CommandLine,CloudOut);

// Copy the output files back to the local computer:
ClusterResult rg = clust.get(CloudWorkingDir, LocalWorkingDir);

// Terminate the cloud cluster:
ClusterResult rt = clust.terminate();
```

End User's view: