

# Supercomputing in Plain English

## Part IX:

### High Throughput Computing

**Henry Neeman, Director**

**OU Supercomputing Center for Education & Research**  
**University of Oklahoma Information Technology**

**Tuesday April 21 2009**





# This is an experiment!

It's the nature of these kinds of videoconferences that  
**FAILURES ARE GUARANTEED TO HAPPEN!**  
**NO PROMISES!**

So, please bear with us. Hopefully everything will work out well enough.

If you lose your connection, you can retry the same kind of connection, or try connecting another way.

Remember, if all else fails, you always have the toll free phone bridge to fall back on.





# Access Grid

This week's Access Grid (AG) venue: Verlet.  
If you aren't sure whether you have AG, you probably don't.

Tue Apr 21	Verlet
Tue Apr 28	Cactus
Tue May 5	Titan

Many thanks to  
John Chapman of  
U Arkansas for  
setting these up  
for us.





## H.323 (Polycom etc)

If you want to use H.323 videoconferencing – for example, Polycom – then dial

**69.77.7.203##12345**

any time after 2:00pm. Please connect early, at least today.

For assistance, contact Andy Fleming of KanREN/Kan-ed ([afleming@kanren.net](mailto:afleming@kanren.net) or 785-230-2513).

KanREN/Kan-ed's H.323 system can handle up to 40 simultaneous H.323 connections. If you cannot connect, it may be that all 40 are already in use.

Many thanks to Andy and KanREN/Kan-ed for providing H.323 access.





# iLinc

We have unlimited simultaneous iLinc connections available.

If you're already on the SiPE e-mail list, then you should already have an e-mail about iLinc. Your personal URL will always be the same.

If you want to use iLinc, please follow the directions in the iLinc e-mail.

For iLinc, you **MUST** use either Windows (XP strongly preferred) or MacOS X with Internet Explorer.

To use iLinc, you'll need to download a client program to your PC. It's free, and setup should take only a few minutes.

Many thanks to Katherine Kantardjieff of California State U Fullerton for providing the iLinc licenses.





# QuickTime Broadcaster

If you cannot connect via the Access Grid, H.323 or iLinc, then you can connect via QuickTime:

**`rtsp://129.15.254.141/test_hpc09.sdp`**

We recommend using QuickTime Player for this, because we've tested it successfully.

We recommend upgrading to the latest version at:

**<http://www.apple.com/quicktime/>**

When you run QuickTime Player, traverse the menus

File -> Open URL

Then paste in the rstp URL into the textbox, and click OK.

Many thanks to Kevin Blake of OU for setting up QuickTime Broadcaster for us.





# Phone Bridge

If all else fails, you can call into our toll free phone bridge:

1-866-285-7778, access code 6483137#

Please mute yourself and use the phone to listen.

Don't worry, we'll call out slide numbers as we go.

Please use the phone bridge **ONLY** if you cannot connect any other way: the phone bridge is charged per connection per minute, so our preference is to minimize the number of connections.

Many thanks to Amy Apon and U Arkansas for providing the toll free phone bridge.





# Please Mute Yourself

No matter how you connect, please mute yourself, so that we cannot hear you.

At OU, we will turn off the sound on all conferencing technologies.

That way, we won't have problems with echo cancellation.

Of course, that means we cannot hear questions.

So for questions, you'll need to send some kind of text.

Also, if you're on iLinc: **SIT ON YOUR HANDS!**

**Please DON'T touch ANYTHING!**







# Questions via Text: iLinc or E-mail

Ask questions via text, using one of the following:

- iLinc's text messaging facility;
- e-mail to [sipe2009@gmail.com](mailto:sipe2009@gmail.com).

All questions will be read out loud and then answered out loud.





# Thanks for helping!

- OSCER operations staff (Brandon George, Dave Akin, Brett Zimmerman, Josh Alexander)
- OU Research Campus staff (Patrick Calhoun, Josh Maxey, Gabe Wingfield)
- Kevin Blake, OU IT (videographer)
- Katherine Kantardjieff, CSU Fullerton
- John Chapman and Amy Apon, U Arkansas
- Andy Fleming, KanREN/Kan-ed
- This material is based upon work supported by the National Science Foundation under Grant No. OCI-0636427, “CI-TEAM Demonstration: Cyberinfrastructure Education for Bioinformatics and Beyond.”





# This is an experiment!

It's the nature of these kinds of videoconferences that  
**FAILURES ARE GUARANTEED TO HAPPEN!**  
**NO PROMISES!**

So, please bear with us. Hopefully everything will work out well enough.

If you lose your connection, you can retry the same kind of connection, or try connecting another way.

Remember, if all else fails, you always have the toll free phone bridge to fall back on.





# Supercomputing Exercises

Want to do the “Supercomputing in Plain English” exercises?

- The first several exercises are already posted at:

<http://www.oscer.ou.edu/education.php>

- If you don’t yet have a supercomputer account, you can get a temporary account, just for the “Supercomputing in Plain English” exercises, by sending e-mail to:

[hneeman@ou.edu](mailto:hneeman@ou.edu)

Please note that this account is for doing the **exercises only**, and will be shut down at the end of the series.





# OK Supercomputing Symposium 2009



**2003 Keynote:**  
**Peter Freeman**  
NSF  
Computer &  
Information  
Science &  
Engineering  
Assistant Director



**2004 Keynote:**  
**Sangtae Kim**  
NSF Shared  
Cyberinfrastructure  
Division Director



**2005 Keynote:**  
**Walt Brooks**  
NASA Advanced  
Supercomputing  
Division Director



**2006 Keynote:**  
**Dan Atkins**  
Head of NSF's  
Office of  
Cyber-  
infrastructure



**2007 Keynote:**  
**Jay Boisseau**  
Director  
Texas Advanced  
Computing Center  
U. Texas Austin



**2008 Keynote:**  
**José Muñoz**  
Deputy Office  
Director/ Senior  
Scientific Advisor  
Office of Cyber-  
infrastructure  
National Science  
Foundation



**2009 Keynote:**  
**Ed Seidel**  
Director  
NSF Office of  
Cyber-  
infrastructure

**FREE! Wed Oct 7 2009 @ OU**  
<http://symposium2009.oscer.ou.edu/>

**Parallel Programming Workshop**  
**FREE! Tue Oct 6 2009 @ OU**

**Registration**  
**is OPEN!**

**Sponsored by SC09 Education Program**  
**FREE! Symposium Wed Oct 7 2009 @ OU**

Supercomputing in Plain English: High Throughput Computing  
Tuesday April 21 2009





# SC09 Summer Workshops

This coming summer, the SC09 Education Program, part of the SC09 (Supercomputing 2009) conference, is planning to hold two weeklong supercomputing-related workshops in Oklahoma, for **FREE** (except you pay your own transport):

- **At OSU Sun May 17 – the May 23:**  
**FREE** Computational Chemistry for Chemistry Educators (2010 TENTATIVE: Computational Biology)
- **At OU Sun Aug 9 – Sat Aug 15:**  
**FREE** Parallel Programming & Cluster Computing

We'll alert everyone when the details have been ironed out and the registration webpage opens.

Please note that you must **apply** for a seat, and acceptance **CANNOT** be guaranteed.





# SC09 Summer Workshops

1. May 17-23: Oklahoma State U: Computational Chemistry
2. May 25-30: Calvin Coll (MI): Intro to Computational Thinking
3. June 7-13: U Cal Merced: Computational Biology
4. June 7-13: Kean U (NJ): Parallel Progrmg & Cluster Comp
5. June 14-20: Widener U (PA): Computational Physics
6. July 5-11: Atlanta U Ctr: Intro to Computational Thinking
7. July 5-11: Louisiana State U: Parallel Progrmg & Cluster Comp
8. July 12-18: U Florida: Computational Thinking Grades 6-12
9. July 12-18: Ohio Supercomp Ctr: Computational Engineering
10. Aug 2- 8: U Arkansas: Intro to Computational Thinking
11. Aug 9-15: U Oklahoma: Parallel Progrmg & Cluster Comp





# Outline

- What is High Throughput Computing?
- Tightly Coupled vs Loosely Coupled
- What is Opportunistic Computing?
- Condor
- Grid Computing
- OU's NSF CI-TEAM Project (a word from our sponsors)





# What is High Throughput Computing?





# High Throughput Computing

**High Throughput Computing** (HTC) means getting lots of work done per large time unit (for example, jobs per month).

This is different from **High Performance Computing** (HPC), which means getting **a particular job** done in less time (for example, calculations per second).





# Throughput vs Performance

- **Throughput** is a side effect of how much time your job takes from when you first submit it until it completes.
- **Performance** is the factor that controls how much time your jobs takes from when it first starts running until it completes.
- Example:
  - You submit a big job at 1:00am on January 1.
  - It sits in the queue for a while.
  - It starts running at 5:00pm on January 2.
  - It finishes running at 6:00pm on January 2.
  - Its performance is fast; its throughput is slow.



# High Throughput on a Cluster?

Is it possible to get high throughput on a cluster?

Sure – it just has to be a cluster that no one else is trying to use!

Normally, a cluster that is shared by many users is fully loaded with jobs all the time. So your throughput depends on when you submit your jobs, and even how many jobs you submit at a time.

Depending on a variety of factors, a job you submit may wait in the queue for anywhere from seconds to days.



# Tightly Coupled vs Loosely Coupled





# Tightly Coupled vs Loosely Coupled

- **Tightly coupled** means that all of the parallel tasks have to advance forward in lockstep, so they have to communicate frequently.
- **Loosely coupled** means that the parallel tasks can largely or completely ignore each other (little or no communication), and they can advance at different rates.





# Tightly Coupled Example


Consider weather forecasting.

You take your simulation domain – for example, the continental United States – split it up into chunks, and give each chunk to an MPI process.

But, the weather in northern Oklahoma affects the weather in southern Kansas.

So, every single timestep, the process that contains northern Oklahoma has to communicate with the process that contains southern Kansas, so that the interface between the processes has the same weather at the same time.

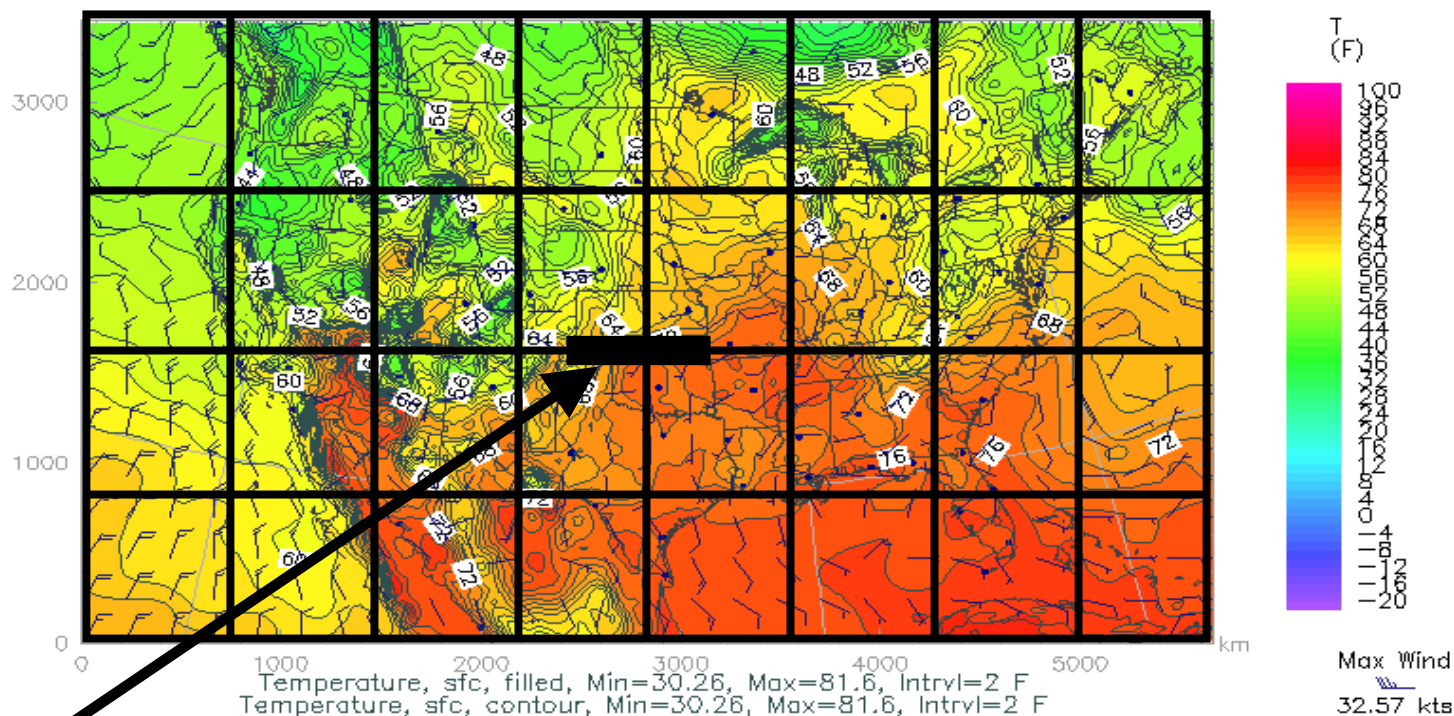




# Tightly Coupled Example

Thu, 25 May 2006, 8 am CDT (13Z)

Surface Temperature



OK/KS  
boundary

<http://www.caps.ou.edu/wx/p/r/conus/fcst/>

CAPS/OU Experimental ADAS Anlys

CONUS, 210x128x50, dx=27 km

05/25/06 08:45 CDT



Supercomputing in Plain English: High Throughput Computing  
Tuesday April 21 2009





# Loosely Coupled Example

An application is known as *embarrassingly parallel*, or *loosely coupled*, if its parallel implementation:

1. can straightforwardly be broken up into roughly equal amounts of work per processor, **AND**
2. has minimal parallel overhead (for example, communication among processors).

We love embarrassingly parallel applications, because they get **near-perfect parallel speedup**, sometimes with only modest programming effort.



# Monte Carlo Methods

Monte Carlo is a city in the tiny European country Monaco.

People gamble there; that is, they play games of chance, which involve randomness.

*Monte Carlo methods* are ways of simulating (or otherwise calculating) physical phenomena based on randomness.


Monte Carlo simulations typically are embarrassingly parallel.





# Monte Carlo Methods: Example

Suppose you have some physical phenomenon. For example, consider High Energy Physics, in which we bang tiny particles together at incredibly high speeds.



We want to know, say, the average properties of this phenomenon.


There are infinitely many ways that two particles can be banged together.

So, we can't possibly simulate all of them.



# Monte Carlo Methods: Example

Suppose you have some physical phenomenon. For example, consider High Energy Physics, in which we bang tiny particles together at incredibly high speeds.



We want to know, say, the average properties of this phenomenon.

There are infinitely many ways that two particles can be banged together.


So, we can't possibly simulate all of them.

Instead, we can randomly choose a finite subset of these infinitely many ways and simulate only the subset.



# Monte Carlo Methods: Example

Suppose you have some physical phenomenon. For example, consider High Energy Physics, in which we bang tiny particles together at incredibly high speeds.



We want to know, say, the average properties of this phenomenon.

There are infinitely many ways that two particles can be banged together.

So, we can't possibly simulate all of them.

The average of this subset will be close to the actual average.



# Monte Carlo Methods

In a Monte Carlo method, you randomly generate a large number of example cases (*realizations*) of a phenomenon, and then take the average of the properties of these realizations.

When the realizations' average converges (that is, doesn't change substantially if new realizations are generated), then the Monte Carlo simulation stops.

This can also be implemented by picking a high enough number of realizations to be sure, mathematically, of convergence.





# MC: Embarrassingly Parallel

Monte Carlo simulations are embarrassingly parallel, because each realization is completely independent of all of the other realizations.

That is, if you're going to run a million realizations, then:

1. you can straightforwardly break up into roughly  $1M / N_p$  chunks of realizations, one chunk for each of the  $N_p$  processes, **AND**
2. the only parallel overhead (for example, communication) comes from tracking the average properties, which doesn't have to happen very often.



# Serial Monte Carlo

Suppose you have an existing serial Monte Carlo simulation:

```
PROGRAM monte_carlo
  CALL read_input(...)
  DO realization = 1, number_of_realizations
    CALL generate_random_realization(...)
    CALL calculate_properties(...)
  END DO
  CALL calculate_average(...)
END PROGRAM monte_carlo
```

How would you parallelize this?





# Parallel Monte Carlo: MPI

```
PROGRAM monte_carlo_mpi
  [MPI startup]
  IF (my_rank == server_rank) THEN
    CALL read_input(...)
  END IF
  CALL MPI_Bcast(...)
  number_of_realizations_per_process = &
& number_of_realizations / number_of_processes
  DO realization = 1, number_of_realizations_per_process
    CALL generate_random_realization(...)
    CALL calculate_realization_properties(...)
    CALL calculate_local_running_average(...)
  END DO
  IF (my_rank == server_rank) THEN
    [collect properties]
  ELSE
    [send properties]
  END IF
  CALL calculate_global_average_from_local_averages(...)
  CALL output_overall_average(...)
[MPI shutdown]
END PROGRAM monte_carlo_mpi
```





# Parallel Monte Carlo: HTC

Suppose you have an existing serial Monte Carlo simulation:

```
PROGRAM monte_carlo
  CALL read_input(...)
  number_of_realizations_per_job = &
&    number_of_realizations / number_of_jobs
  DO realization = 1, number_of_realizations_per_job
    CALL generate_random_realization(...)
    CALL calculate_properties(...)
  END DO
  CALL calculate_average_for_this_job(...)
  CALL output_average_for_this_job(...)
END PROGRAM monte_carlo
```

To parallelize this for HTC, simply submit **number\_of\_jobs** jobs, and then at the very end run a little program to calculate the overall average.




# What is Opportunistic Computing?





# Desktop PCs Are Idle Half the Day



Desktop PCs tend to be active during the workday.



But at night, during most of the year, they're idle. So we're only getting half their value (or less).



# Supercomputing at Night

A particular institution – say, OU – has lots of desktop PCs that are **idle during the evening and during intersessions**.

Wouldn't it be great to put them to work on something **useful** to our institution?

That is: What if they could pretend to be a big supercomputer **at night**, when they'd **otherwise be idle anyway**?

This is sometimes known as **opportunistic computing**:

When a desktop PC is otherwise idle, you have an opportunity to do number crunching on it.





# Supercomputing at Night Example

**SETI** – the Search for Extra-Terrestrial Intelligence – is looking for evidence of green bug-eyed monsters on other planets, by mining radio telescope data.

**SETI@home** runs number crunching software as a screensaver on idle PCs around the world (2+ million PCs in 252 countries):

<http://setiathome.berkeley.edu/>



There are **many similar projects**:

- folding@home (protein folding)
- climateprediction.net
- Einstein@Home (Laser Interferometer Gravitational wave Observatory)
- Cosmology@home
- ...





# BOINC

The projects listed on the previous page use a software package named BOINC (**B**erkeley **O**pen **I**nfrastructure for **N**etwork **C**omputing), developed at the University of California, Berkeley:

<http://boinc.berkeley.edu/>



To use BOINC, you have to insert calls to various BOINC routines into your code. It looks a bit similar to MPI:

```
int main ()
{ /* main */
    ...
    boinc_init();
    ...
    boinc_finish(...);
} /* main */
```





# Condor







# Condor is Like BOINC

- Condor steals computing time on existing desktop PCs when they're idle.
- Condor runs in background when no one is sitting at the desk.
- Condor allows an institution to get much more value out of the hardware that's already purchased, because there's little or no idle time on that hardware – all of the idle time is used for number crunching.





# Condor is Different from BOINC

- To use Condor, **you don't need to rewrite your software** to add calls to special routines; in BOINC, you do.
- Condor **works great under Unix/Linux**, but less well under Windows or MacOS (more on this presently); BOINC works well under all of them.
- It's **non-trivial to install Condor** on your own personal desktop PC; it's straightforward to install a BOINC application such as SETI@home.





# Useful Features of Condor

- **Opportunistic** computing: Condor steals time on existing desktop PCs when they're otherwise not in use.
- Condor **doesn't require any changes to the software.**
- Condor can **automatically checkpoint** a running job: Every so often, Condor saves to disk the state of the job (the values of all the job's variables, plus where the job is in the program).
- Therefore, Condor can **preempt** running jobs if more important jobs come along, or if someone sits down at the desktop PC.
- Likewise, Condor can **migrate** running jobs to other PCs, if someone sits at the PC or if the PC crashes.
- And, Condor can do all of its **I/O over the network**, so that the job on the desktop PC doesn't consume the desktop PCs local disk.



# Condor Pool @ OU


OU IT has deployed a large Condor pool  
(773 desktop PCs in dozens of labs around campus).

OU's Condor pool provides a huge amount of  
computing power – more than OSCER's big  
cluster:

- if OU were a state, we'd be the 16th largest  
state in the US;
- if OU were a country, we'd be the 13th largest  
country in the world.

The hardware and software cost is zero, and the  
labor cost is modest.

Also, we've been seeing empirically that lab PCs  
are available for Condor jobs about 80% of the time.





# Condor Limitations

- The Unix/Linux version has more features than Windows or MacOS, which are referred to as “clipped.”
- Your code shouldn't be parallel to do opportunistic computing (MPI requires a fixed set of resources throughout the entire run), and it shouldn't try to do any funky communication (for example, opening sockets).
- For a Red Hat Linux Condor pool, you have to be able to compile your code with gcc, g++, g77 or NAG f95 (which is a Fortran90-to-C translator that then calls gcc).
- Also, depending on the PCs that have Condor on them, you may have limitations on, for example, how big your jobs' RAM footprint can be.





# Running a Condor Job

Running a job on Condor pool is a lot like running a job on a cluster:

1. You compile your code using the compilers appropriate for that resource.
2. You submit a batch script to the Condor system, which decides when and where your job runs, magically and invisibly.





# Sample Condor Batch Script

```
Universe      = standard
Executable    = /home/hneeman/NBody/nbody_compiled_for_condor
Notification  = Error
Notify_User   = hneeman@ou.edu
Arguments     = 1000 100
Input         = /home/hneeman/NBody/nbody_input.txt
Output        = nbody_$(Cluster)_$(Process)_out.txt
Error         = nbody_$(Cluster)_$(Process)_err.txt
Log           = nbody_$(Cluster)_$(Process)_log.txt
InitialDir    = /home/hneeman/NBody/Run001
Queue
```

The batch submission command is **condor\_submit**, used like so:

```
condor_submit nbody.condor
```





# Linux Condor on Windows PCs?

If OU's Condor pool uses Linux, how can it be installed in OU IT PC labs? Don't those run Windows?

**Yes.**

Our solution is to run Linux inside Windows, using a piece of software named coLinux (“Cooperative Linux”):

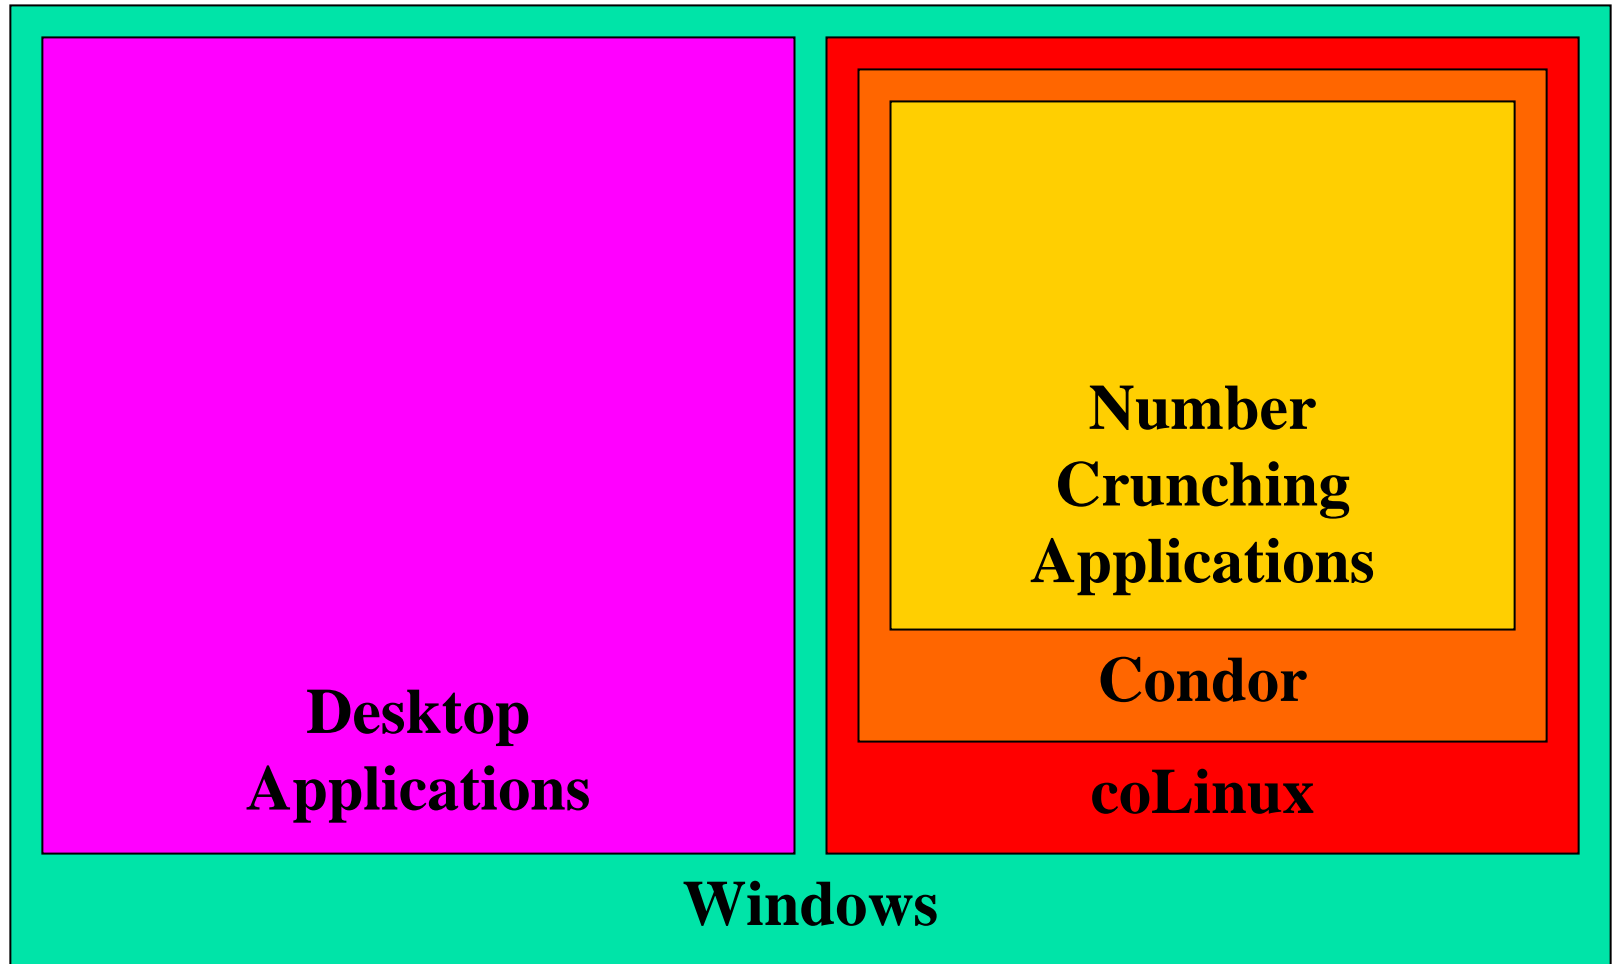
<http://www.colinux.org/>







# Condor inside Linux inside Windows





# Advantages of Linux inside Windows

- Condor is full featured rather than clipped.
- Desktop users have a full Windows experience, without even being aware that coLinux exists.
- A little kludge helps Condor watch the keyboard, mouse and CPU level of Windows, so that Condor jobs don't run when the PC is otherwise in use.

Want to try it yourself?

[http://www.oscer.ou.edu/CondorInstall/condor\\_colinux\\_howto.php](http://www.oscer.ou.edu/CondorInstall/condor_colinux_howto.php)





# Grid Computing

---



# What is Grid Computing?

The term *grid computing* is poorly defined, but the best definition I've seen so far is:

“a distributed, heterogeneous operating system.”

A *grid* can consist of:

- compute resources;
- storage resources;
- networks;
- data collections;
- shared instruments;
- sensor networks;
- and so much more ....





# Grid Computing is Like and Unlike ...

IBM's website has a very good description of grid computing:

- *“Like the Web, grid computing keeps complexity hidden: multiple users enjoy a single, unified experience.*
- *“Unlike the Web, which mainly enables communication, grid computing enables full collaboration toward common ... goals.*
- *“Like peer-to-peer, grid computing allows users to share files.*
- *“Unlike peer-to-peer, grid computing allows many-to-many sharing – not only files but other resources as well.*
- *“Like clusters and distributed computing, grids bring computing resources together.*
- *“Unlike clusters and distributed computing, which need physical proximity and operating homogeneity, grids can be geographically distributed and heterogeneous.*
- *“Like virtualization technologies, grid computing enables the virtualization of IT resources.*
- *“Unlike virtualization technologies, which virtualize a single system, grid computing enables the virtualization of vast and disparate IT resources.”*

[http://www.thocp.net/hardware/grid\\_computers.htm](http://www.thocp.net/hardware/grid_computers.htm)






# Condor is Grid Computing

Condor creates a grid out of disparate desktop PCs.

(Actually, they don't have to be desktop PCs; they don't even have to be PCs. You can use Condor to schedule a cluster, or even on a big iron supercomputer.)

From a user's perspective, all of the PCs are essentially invisible; the user just knows how to submit a job, and everything happens magically and invisibly, and at some point the job is done and a result appears.





# OU's NSF CI-TEAM Project

---



# NSF CI-TEAM Grant

“Cyberinfrastructure Education for Bioinformatics and Beyond” (\$249,976, 12/01/2006 – 11/30/2008, no cost extension to 11/30/2009)

OSCER received a grant from the National Science Foundation’s Cyberinfrastructure Training, Education, Advancement, and Mentoring for Our 21st Century Workforce (CI-TEAM) program.







# OU's NSF CI-TEAM Grant

“Cyberinfrastructure Education for Bioinformatics and Beyond” (\$249,976)

Objectives:

- Provide Condor resources to the national community.
- Teach users to use Condor.
- Teach sysadmins to deploy and administer Condor.
- Teach supercomputing to everyone!

[You can join!](#)



Supercomputing in Plain English: High Throughput Computing  
Tuesday April 21 2009



# OU's NSF CI-TEAM Grant

## Participants at OU

(29 faculty/staff in 16 depts)

- Information Technology
  - OSCER: Neeman (PI)
- College of Arts & Sciences
  - Botany & Microbiology: Conway, Wren
  - Chemistry & Biochemistry: Roe (Co-PI), Wheeler
  - Mathematics: White
  - Physics & Astronomy: Kao, Severini (Co-PI), Skubic, Strauss
  - Zoology: Ray
- College of Earth & Energy
  - Sarkeys Energy Center: Chesnokov
- College of Engineering
  - Aerospace & Mechanical Engr: Striz
  - Chemical, Biological & Materials Engr: Papavassiliou
  - Civil Engr & Environmental Science: Vieux
  - Computer Science: Dhall, Fagg, Hougen, Lakshmivarahan, McGovern, Radhakrishnan
  - Electrical & Computer Engr: Cruz, Todd, Yeary, Yu
  - Industrial Engr: Trafalis
- Health Sciences Center
  - Biochemistry & Molecular Biology: Zlotnick
  - Radiological Sciences: Wu (Co-PI)
  - Surgery: Gusev

## Participants at other institutions

(46 faculty/staff at 30 institutions in 18 states)


1. California State U Pomona (masters-granting, minority serving): Lee
2. Colorado State U: Kalkhan
3. Contra Costa College (CA, 2-year, minority serving): Murphy
4. Delaware State U (masters, EPSCoR): Hubsch, Mulik, Multnovic, Rasamny
5. Earlham College (IN, bachelors): Peck
6. East Central U (OK, masters, EPSCoR): Ferdinand, Myers
7. Emporia State U (KS, masters-granting, EPSCoR): Pheatt, Ballester
8. Harvard U (MA): Altman
9. Kansas State U (EPSCoR): Andresen, Monaco
10. Langston U (OK, masters, minority serving, EPSCoR): Depona, Snow, Tadesse
11. Longwood U (VA, masters): Talaiver
12. Marshall U (WV, masters, EPSCoR): Richards
13. Navajo Technical College (NM, 2-year, tribal, EPSCoR): Ribble
14. Oklahoma Baptist U (bachelors, EPSCoR): Chen, Jett, Jordan
15. Oklahoma Medical Research Foundation (EPSCoR): Wren
16. Oklahoma School of Science & Mathematics (high school, EPSCoR): Samadzadeh
17. Riverside Community College (CA, 2-year): Smith
18. St. Cloud State University (MN, masters): Herath
19. St. Gregory's U (OK, 4-year, EPSCoR): Meyer
20. Southwestern Oklahoma State U (masters, EPSCoR, tribal): Linder, Moseley
21. Syracuse U (NY): Chen
22. Texas A&M U-Corpus Christi (masters): Scherger
23. U Arkansas Fayetteville (EPSCoR): Apon
24. U Arkansas Little Rock (masters, EPSCoR): Jennings, Ramaswamy
25. U Central Oklahoma (masters-granting, EPSCoR): Lemley, Wilson
26. U Illinois Urbana-Champaign: Wang
27. U Kansas (EPSCoR): Bishop, Cheung, Harris, Ryan
28. U Nebraska-Lincoln (EPSCoR): Swanson
29. U North Dakota (EPSCoR): Bergstrom
30. U Northern Iowa (masters-granting): Gray



Supercomputing in Plain English: High Throughput Computing

Tuesday April 21 2009

# NSF CI-TEAM Participants





# NSF CI-TEAM Grant

“Cyberinfrastructure Education for Bioinformatics and Beyond” (\$249,976)

OSCER is providing “Supercomputing in Plain English” workshops via videoconferencing starting in Fall 2007.

434 registrants at 70+ institutions, agencies, companies and non-governmental organizations in 29 US states and Puerto Rico as well as Mexico, Argentina, Switzerland and India, via:

- Access Grid
- H.323 (for example, Polycom)
- iLinc
- QuickTime
- Phone bridge (land line)



Supercomputing in Plain English: High Throughput Computing  
Tuesday April 21 2009



# SiPE Workshop Participants Fall 2009



Supercomputing in Plain English: High Throughput Computing  
Tuesday April 21 2009



# NSF CI-TEAM Grant

“Cyberinfrastructure Education for Bioinformatics and Beyond” (\$249,976)

OSCER has produced software for installing Linux-enabled Condor inside a Windows PC.

**INTERESTED?** Contact Henry ([hneeman@ou.edu](mailto:hneeman@ou.edu))





# NSF CI-TEAM Grant

“Cyberinfrastructure Education for Bioinformatics and Beyond” (\$249,976)

OSCER is providing help on installing Windows as the native host OS, coLinux inside Windows, Linux inside coLinux and Condor inside Linux.

**INTERESTED?** Contact Henry ([hneeman@ou.edu](mailto:hneeman@ou.edu))







# OK Supercomputing Symposium 2009



**2003 Keynote:**  
**Peter Freeman**  
NSF  
Computer &  
Information  
Science &  
Engineering  
Assistant Director



**2004 Keynote:**  
**Sangtae Kim**  
NSF Shared  
Cyberinfrastructure  
Division Director



**2005 Keynote:**  
**Walt Brooks**  
NASA Advanced  
Supercomputing  
Division Director



**2006 Keynote:**  
**Dan Atkins**  
Head of NSF's  
Office of  
Cyber-  
infrastructure



**2007 Keynote:**  
**Jay Boisseau**  
Director  
Texas Advanced  
Computing Center  
U. Texas Austin



**2008 Keynote:**  
**José Muñoz**  
Deputy Office  
Director/ Senior  
Scientific Advisor  
Office of Cyber-  
infrastructure  
National Science  
Foundation



**2009 Keynote:**  
**Ed Seidel**  
Director  
NSF Office of  
Cyber-  
infrastructure

**FREE! Wed Oct 7 2009 @ OU**  
<http://symposium2009.oscer.ou.edu/>

**Parallel Programming Workshop**  
**FREE! Tue Oct 6 2009 @ OU**

**Registration**  
**is OPEN!**

**Sponsored by SC09 Education Program**  
**FREE! Symposium Wed Oct 7 2009 @ OU**

Supercomputing in Plain English: High Throughput Computing  
Tuesday April 21 2009







# SC09 Summer Workshops

This coming summer, the SC09 Education Program, part of the SC09 (Supercomputing 2009) conference, is planning to hold two weeklong supercomputing-related workshops in Oklahoma, for **FREE** (except you pay your own transport):

- **At OSU Sun May 17 – the May 23:**  
**FREE** Computational Chemistry for Chemistry Educators (2010 TENTATIVE: Computational Biology)
- **At OU Sun Aug 9 – Sat Aug 15:**  
**FREE** Parallel Programming & Cluster Computing

We'll alert everyone when the details have been ironed out and the registration webpage opens.

Please note that you must **apply** for a seat, and acceptance **CANNOT** be guaranteed.





# SC09 Summer Workshops

1. May 17-23: Oklahoma State U: Computational Chemistry
2. May 25-30: Calvin Coll (MI): Intro to Computational Thinking
3. June 7-13: U Cal Merced: Computational Biology
4. June 7-13: Kean U (NJ): Parallel, Distributed & Grid
5. June 14-20: Widener U (PA): Computational Physics
6. July 5-11: Atlanta U Ctr: Intro to Computational Thinking
7. July 5-11: Louisiana State U: Parallel, Distributed & Grid
8. July 12-18: U Florida: Computational Thinking Pre-college
9. July 12-18: Ohio Supercomp Ctr: Computational Engineering
10. Aug 2- 8: U Arkansas: Intro to Computational Thinking
11. Aug 9-15: U Oklahoma: Parallel, Distributed & Grid





# To Learn More Supercomputing

<http://www.oscer.ou.edu/education.php>



Supercomputing in Plain English: High Throughput Computing  
Tuesday April 21 2009



# Thanks for helping!

- OSCER operations staff (Brandon George, Dave Akin, Brett Zimmerman, Josh Alexander)
- OU Research Campus staff (Patrick Calhoun, Josh Maxey, Gabe Wingfield)
- Kevin Blake, OU IT (videographer)
- Katherine Kantardjieff, CSU Fullerton
- John Chapman and Amy Apon, U Arkansas
- Andy Fleming, KanREN/Kan-ed
- This material is based upon work supported by the National Science Foundation under Grant No. OCI-0636427, “CI-TEAM Demonstration: Cyberinfrastructure Education for Bioinformatics and Beyond.”



**Thanks for your  
attention!**



**Questions?**