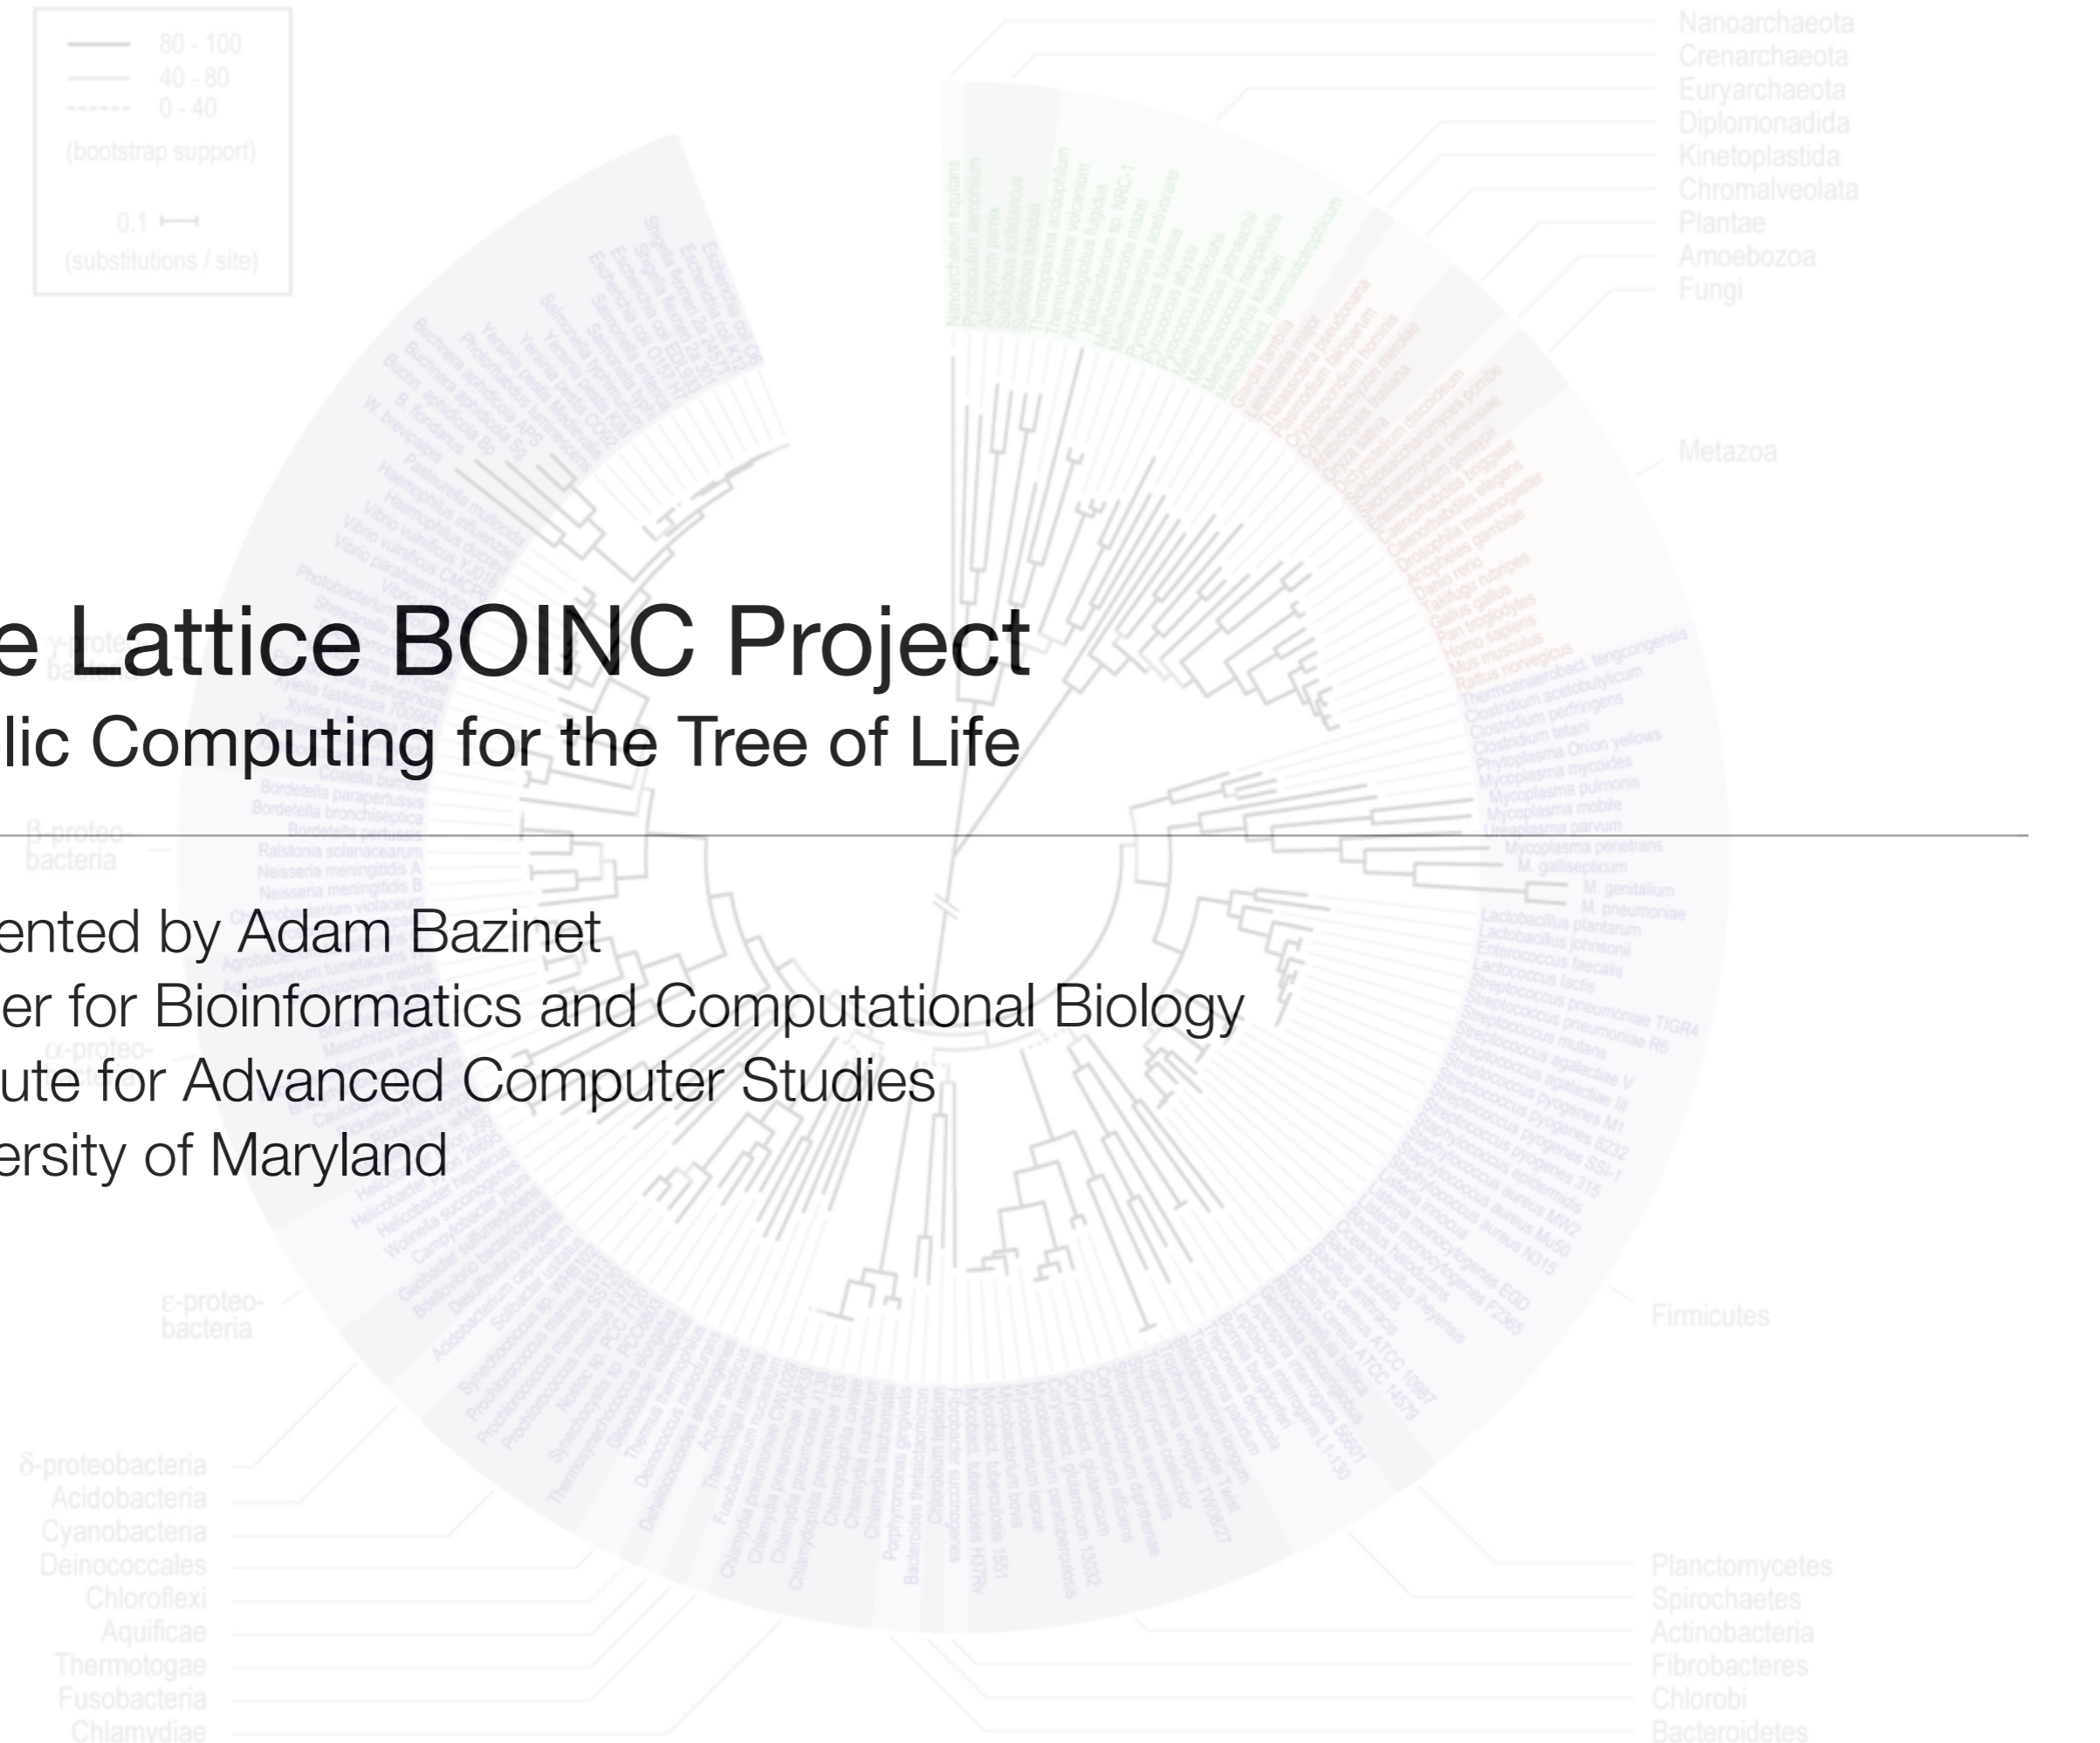


The Lattice BOINC Project


Public Computing for the Tree of Life

Presented by Adam Bazinet
 Center for Bioinformatics and Computational Biology
 Institute for Advanced Computer Studies
 University of Maryland



Parallel computing “universe”

Multi-core



GPU



SMP



Cluster



MPP




Grid computing



Parallel computing “universe”

Shared memory

Multi-core



GPU




SMP



Parallel computing “universe”

Multi-core



GPU



SMP



Cluster



MPP



Grid computing



Parallel computing “universe”

Distributed memory

Cluster



MPP



Grid computing



Parallel computing “universe”

“Distributed computing”

Cluster



MPP




Grid computing



Parallel computing “universe”

Multi-core



GPU



SMP



Cluster



MPP



Grid computing



Parallel computing “universe”

GPU



SMP



Cluster



MPP



Multi-core



Grid computing



Grid computing

BOINC pool



Compute cluster



Condor pool



The Lattice Project: resource characterization

HTC resources

Condor pool



BOINC pool



HTC = high *throughput* computing

HPC resources

Compute cluster



HPC = high *performance* computing

The Lattice Project: resource characterization

Shared

Condor pool



BOINC pool



shared = resources usually have interactive users

Dedicated

Compute cluster



dedicated = resources dedicated to processing

The Lattice Project: resource characterization

Institutional

Condor pool Compute cluster



institutional = resources belong to an institution

Volunteer

BOINC pool



dedicated = resources volunteered by the public

The Lattice Project: resource characterization

Trusted

Condor pool Compute cluster



trusted = results need not be verified

Untrusted

BOINC pool



untrusted = results need verification

The Lattice Project: resource characterization

Reliable

Condor pool Compute cluster



reliable = resources are stable and accounted for

Unreliable

BOINC pool



unreliable = resources are dynamic and unpredictable

Models of Grid computing

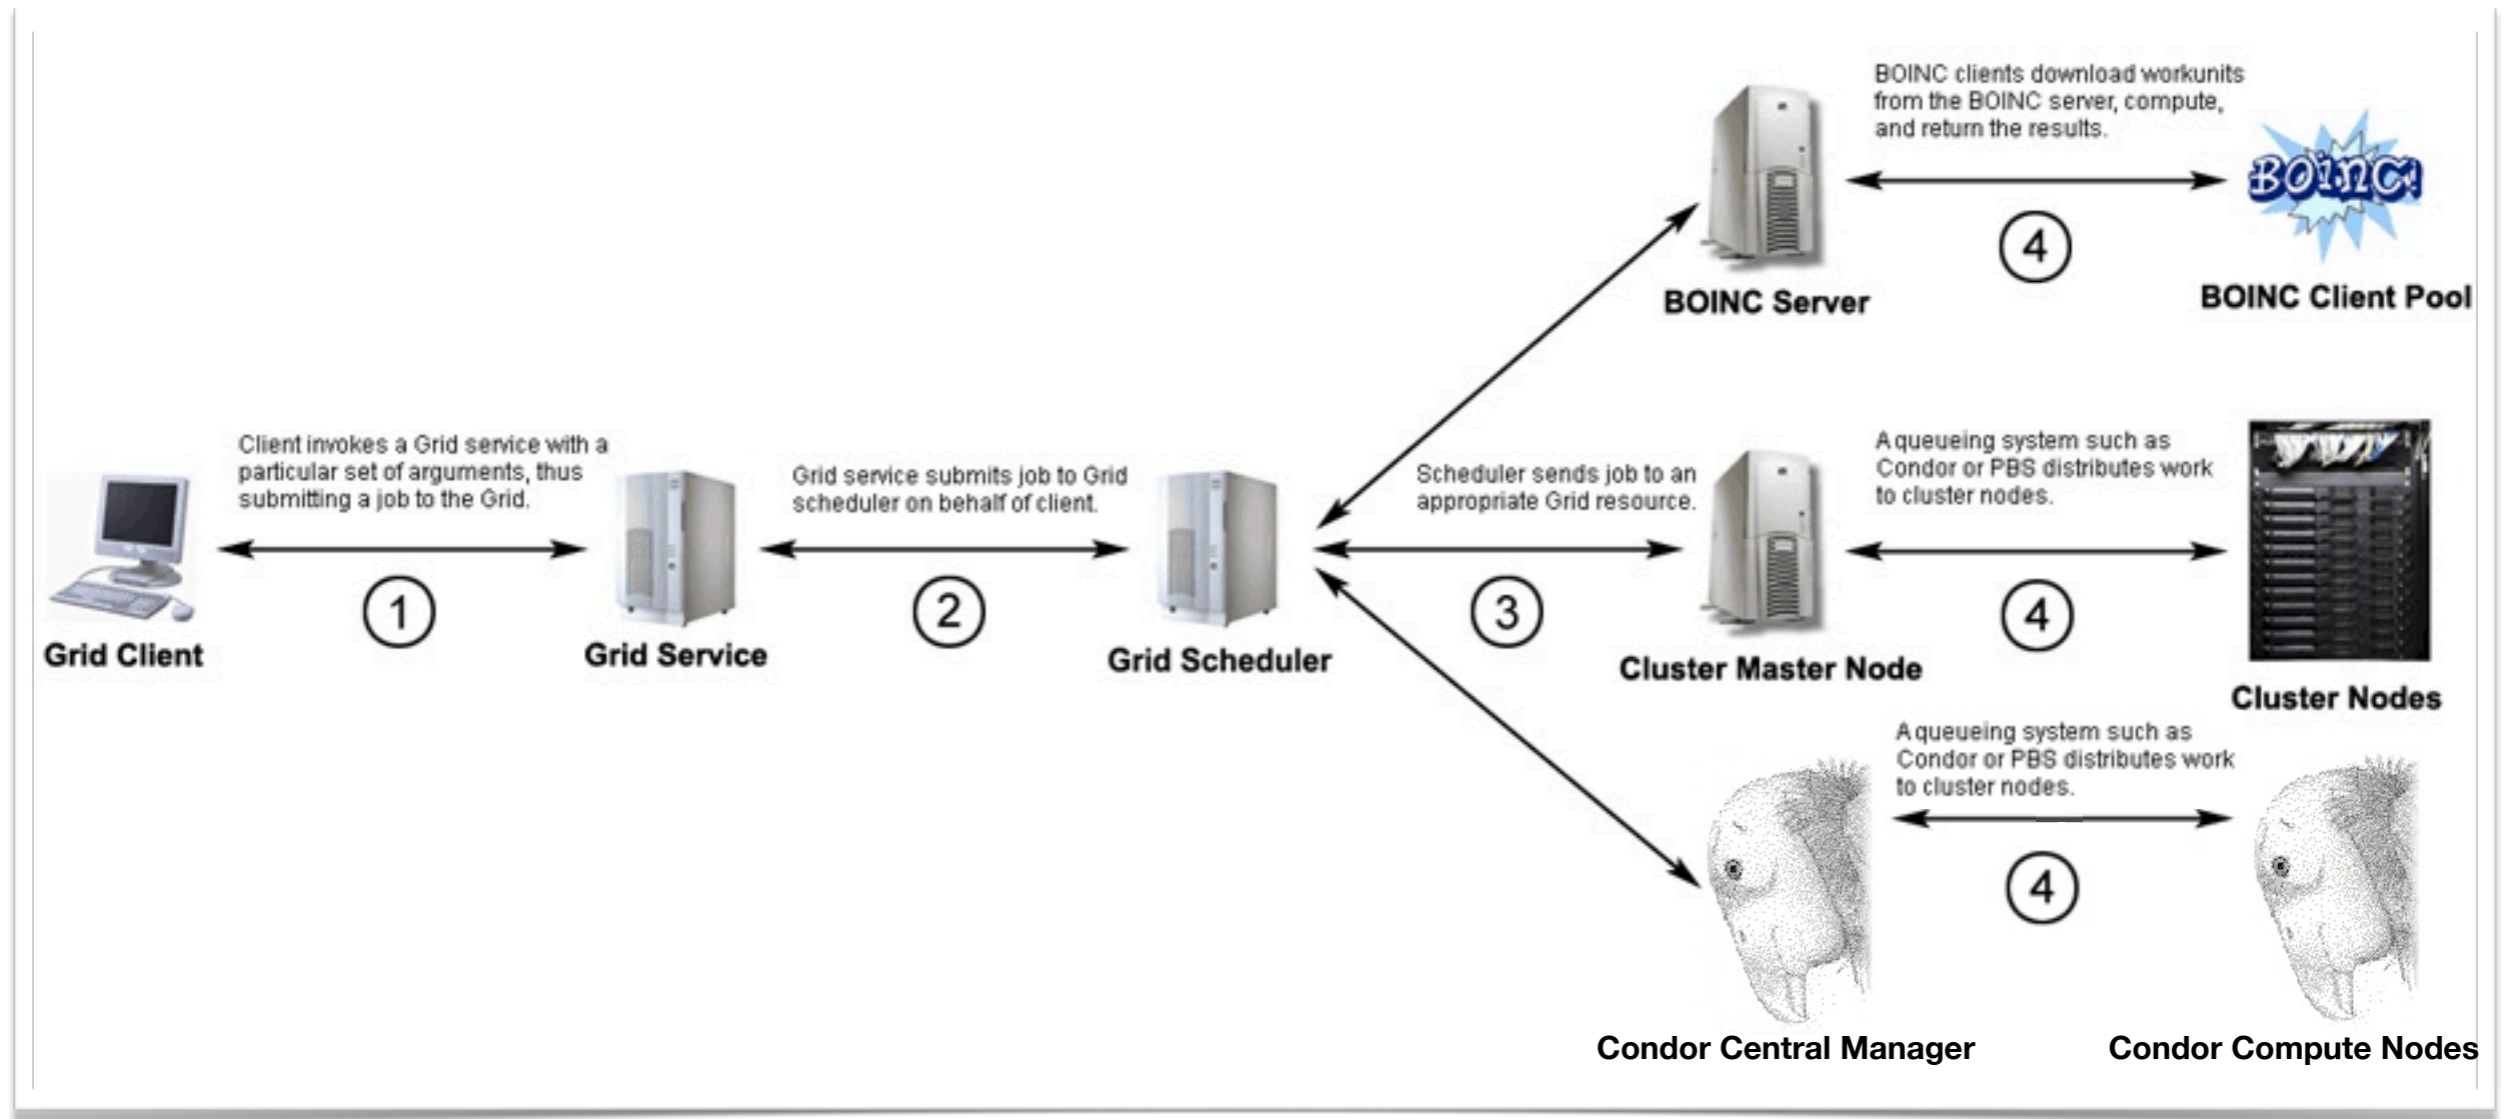


- SERVICE MODEL — a heavyweight, feature-rich model focused on providing access to institutional resources and robust job submission capabilities and security features
 - Well known Service Grids include TeraGrid, Open Science Grid, and EGEE
- DESKTOP MODEL — scavenges cycles from idle desktop computers, which are volunteered by the general public
 - The combined power of hundreds of thousands of desktop computers represents a substantial, readily available resource
 - The most widely used software for tapping this resource is BOINC

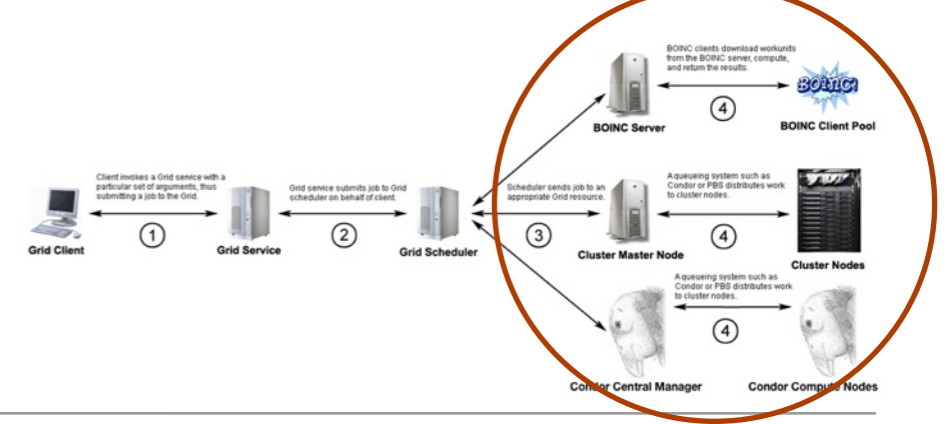


- The first Grid system to effectively combine a Service Grid (using Globus software) and a Desktop Grid (using BOINC software)
- Aimed at sharing computational resources between academic institutions, particularly those in the University System of Maryland
- Focused on enabling large-scale computation, especially for problems in the life sciences
- Development began in 2003 — since then, many different researchers have used the system, racking up over 20,000 CPU years of computation (measured in wall clock time)

The Lattice Project: architecture



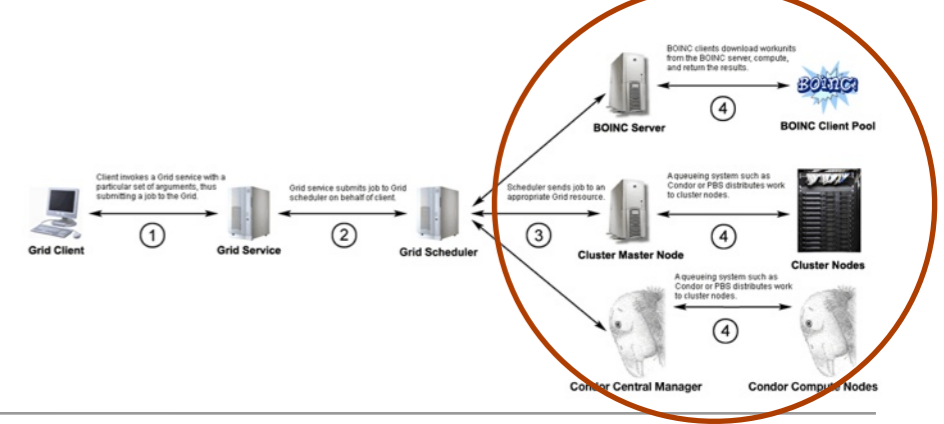
The Lattice Project: resources



- Quick facts about resources:

- We support three major platforms: Linux (both PowerPC and Intel-based), Windows, and Mac (both PowerPC and Intel-based)
- Four different institutions are currently tied in to the Grid: UMCP, Bowie State University, Coppin State University, and the Smithsonian NMNH
- We currently have four Condor pools, four dedicated clusters, and a BOINC project with a steady number of participants
- We currently have a total of 4000-5000 CPUs

The Lattice Project: resources




<http://lattice.umiacs.umd.edu/resources/>

<p>UMIACS Condor Pool 252.26 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>INTEL_LINUX 0 0 490</p> <p>Disk used: 19.7G / 91.2G</p> <p>show/hide condor_status</p>	<p>Coppin Condor Pool 64.21 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>INTEL_WIN 40 0 505</p> <p>Disk used: 4.5G / 13.4G</p> <p>show/hide condor_status</p>	<p>Terpcondor Condor Pool 402.59 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>INTEL_WIN 0 0 185</p> <p>Disk used: 0.2G / 59.8G</p> <p>show/hide condor_status</p>
<p>Topaz 0.00 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>INTEL_OSX 0 0 0</p> <p>Disk used: 158.4G / 0.9T</p>	<p>Lattice on BOINC 17318.11 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>INTEL_LINUX 0 0 125</p> <p>PPC_OSX 0 0 114</p> <p>INTEL_WIN 0 0 1377</p> <p>Disk used: 1.2G / 60.7G</p> <p>Lattice on BOINC web site</p>	<p>Bowie Condor Pool 5.33 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>INTEL_WIN 0 0 0</p> <p>Disk used: 5.2G / 72.5G</p> <p>show/hide condor_status</p>
<p>Xseed 228.19 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>PPC_LINUX 0 0 408</p> <p>Disk used: 41.4G / 1.1T</p> <p>Xseed web site Ganglia</p>	<p>Deepthought 641.87 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>INTEL_LINUX 0 0 112</p> <p>Disk used: 345.6G / 2.5T</p> <p>HPCC web site Ganglia Stats</p>	<p>SEIL 325.29 CPU Years</p> <p>Lattice Jobs</p> <p>Arch. & OS Idle Running Free CPUs</p> <p>INTEL_LINUX 0 0 112</p> <p>Disk used: 67.8G / 122.2G</p> <p>HPCC web site Ganglia Stats</p>
<p>Total Lattice Jobs Idle 40 Running 0</p> <p>Total Free CPUs* Linux 1247 Windows 2067 Mac OS X 114 Grand Total 3428</p> <p>Grid Server Disk Used 1.0G / 74.7G</p>		
<p><small>*NOTE: multi-core machines report each core as a separate CPU.</small></p>		

The Lattice BOINC Project



<http://boinc.umiacs.umd.edu/>



The Lattice Project


Center for Bioinformatics and Computational Biology

[Home](#) [Research](#) [Publications](#) [Personnel](#)

[Home](#)

The Lattice Project

- About Lattice
- Applications
- Create Account
- Message Boards
- Participant Profiles
- Questions & Answers
- Research Projects
- Rules and Policies
- Statistics
- Teams
- Top Computers
- Top Participants
- Top Teams
- Your Account



About The Lattice Project

- [Background and introduction](#)
- [Research projects](#)
- [Applications](#)

Join The Lattice Project

- [Read our rules and policies](#)
- [Getting started](#)
- [Create account](#)
- [Download BOINC](#)
- When prompted, enter <http://boinc.umiacs.umd.edu/>

Returning participants

- [Your account](#)
- [Teams](#)


Community

- [Participant profiles](#)
- [Message boards](#)
- [Questions and answers](#)

Project totals and leader boards

- [Top participants](#)
- [Top computers](#)
- [Top teams](#)
- [Other statistics](#)

User of the Day



Bruno Ramone
Cześć boincomaniacy !

w projektach rozproszonych jestem od 2000 roku, kiedy to zainteresował mnie niezmiernie wygaszacz projektu...

Project Status

BOINC Server	Results	Users
Scheduler: running	Results ready to send: 2	Total Users: 22152
File Server: running	Results in progress: 815	Total Hosts: 38961
Feeder: running	Workunits in database: 26686	


[[Detailed Status](#)] [[Workunit Graveyard](#)]

News

October 9, 2009 - **Server Upgrade**
Just upgraded the server software, so please let me know if you notice anything out of place.

August 1, 2009 - **Workunit Errors**
A number of GARLI jobs were accidentally released with an error in the configuration file that caused the job to terminate incorrectly; these jobs have been removed from the system, and we apologize for this.

[[Old News](#)]

News is available as an [RSS feed](#) .

BOINC as part of a Grid computing system

- The Lattice BOINC Project is part of a Grid computing system that also incorporates traditional computing clusters and Condor pools
- How it works: a job with a generic description is submitted by a Grid user, which is transformed into either a BOINC workunit, a PBS submit file, or a Condor submit file by the Globus GRAM service
- As a consequence of being hooked up to a general-purpose Grid computing system, our project runs many different types of applications
- As a result, we have had to work harder to manage volunteer expectations since we do not have an endless supply of work and the applications (and workunits within an application) are highly variable

The Lattice BOINC Project: research

- Phylogenetic analysis – GARLI



- Protein sequence comparison – HMMPfam

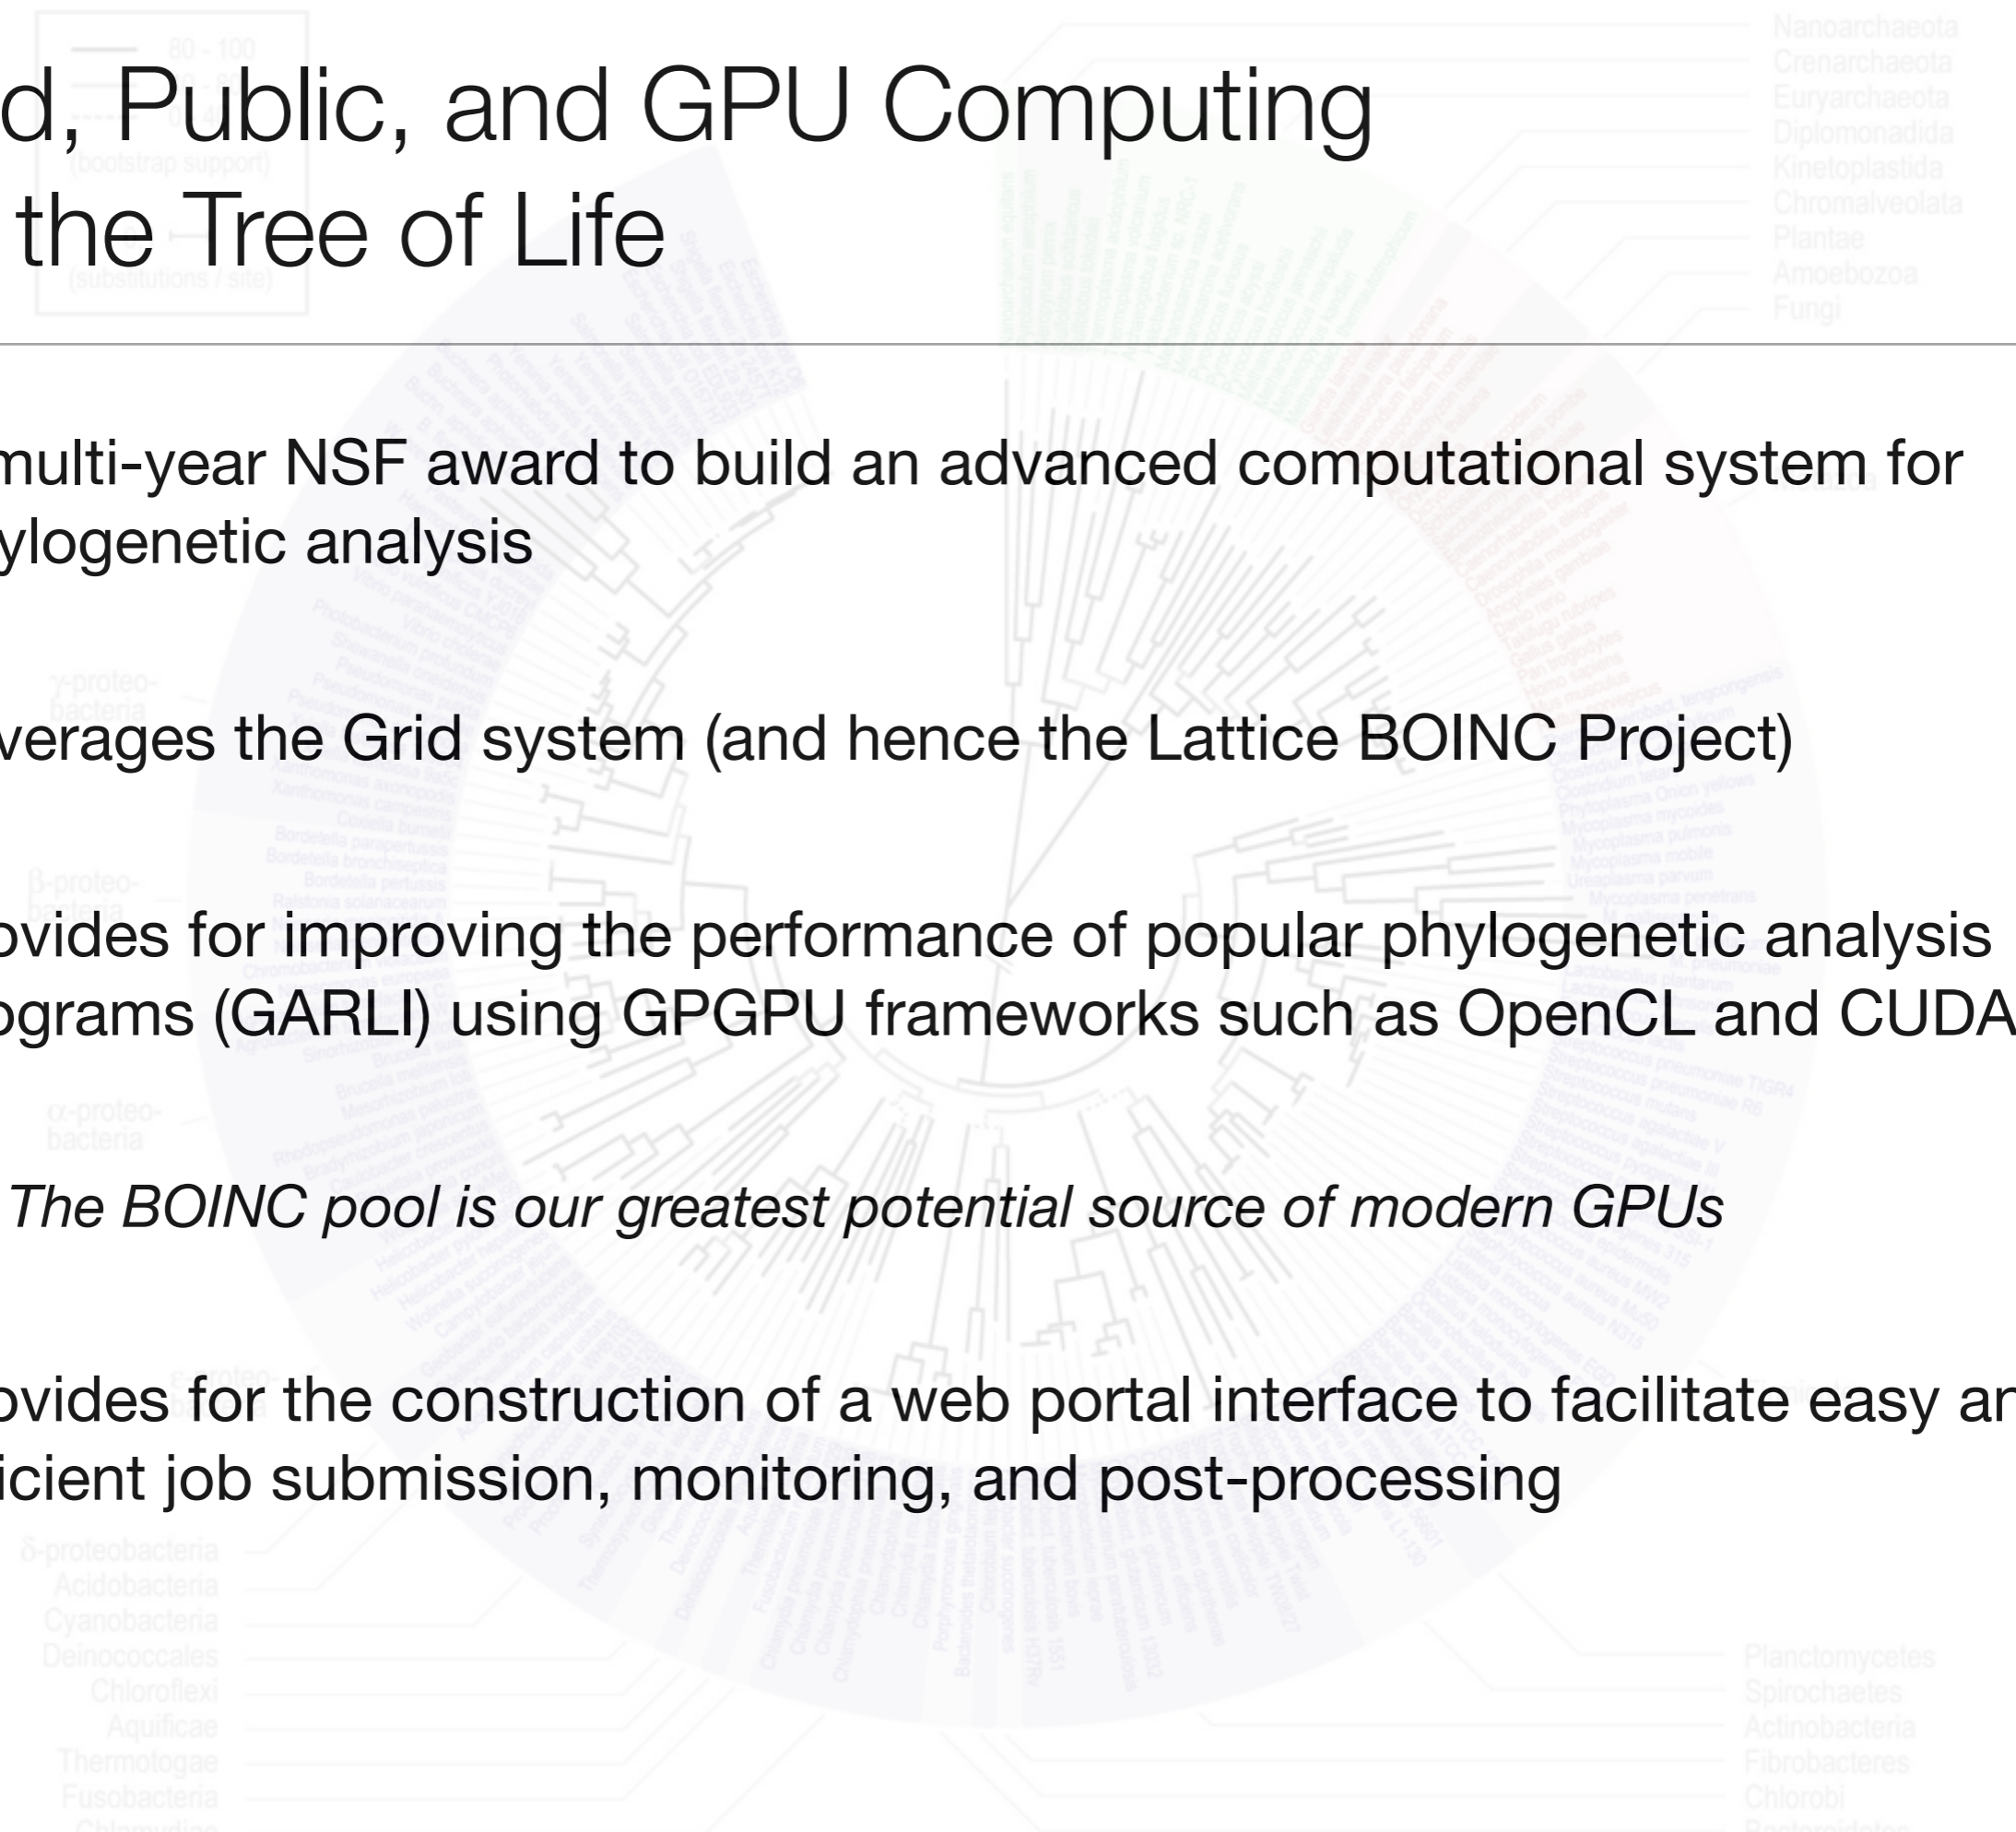


- Conservation network design – MARXAN



Grid, Public, and GPU Computing for the Tree of Life

- A multi-year NSF award to build an advanced computational system for phylogenetic analysis
- Leverages the Grid system (and hence the Lattice BOINC Project)
- Provides for improving the performance of popular phylogenetic analysis programs (GARLI) using GPGPU frameworks such as OpenCL and CUDA
 - *The BOINC pool is our greatest potential source of modern GPUs*
- Provides for the construction of a web portal interface to facilitate easy and efficient job submission, monitoring, and post-processing

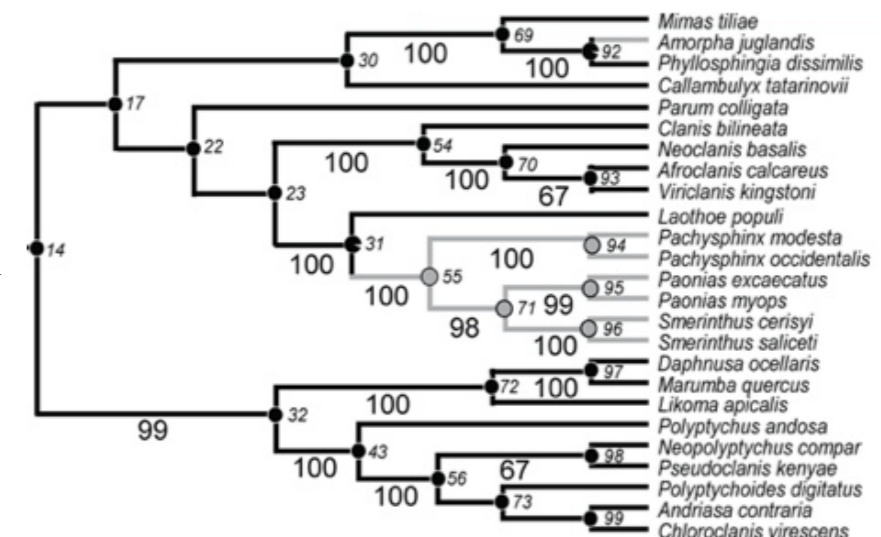




GARLI



- Genetic Algorithm for Rapid Likelihood Inference
- Makes maximum likelihood phylogenetic searches feasible for large datasets
- Avoids calculating the exact “maximized likelihood” for each topology
- Gradually and independently optimizes tree topology, branch lengths, and model parameters

```
ATG---TTC-----ATAATTAACATCTTAATACTA---  
ATG---CTAAACATCCTAATAACTCACCTAATCAACCCT  
ATGACCCTGCCACCCTAACAAACCTTCTAATCATAACC  
ATA---CCC-----ATGGCCAACCTCCTACTCCTC---  
ATC---CTAAACACTCTACTCACCCACCTCATCAACCCC  
ATAGTGTT-----TTTATTAATATCCTAACACTC---  
ATAGTGAC-----TTTATTAATATCCTAACACTC---  
ATG---TTT-----ATAATTAATATCATCTCACTA---  
ATG---TTT-----ATAATTAACATTCTAACACTC---
```



akawahar

My account

Log out

[Home](#) » [GARLI Web Service - create job](#)

GARLI Web Service – create job

GARLI – Genetic Algorithm for Rapid Likelihood Inference

Version 1.0 – Author(s): Derrick J. Zwickl (zwickl@ku.edu) – Category: Phylogenetics

Job Information

Current user:

*** Job name (use only a-z A-Z 0-9 . _):**

General Settings

*** Analysis type:**

*** Number of replicates (1 – 2000):**

Sequence data file *:

akawahar

My account

Log out

[Home](#) » [GARLI Web Service](#) – view job status

GARLI Web Service – view job status

[< back to create job](#)*akawahar*

ID	Job name	Completed	Total	Status	Created	Remove job
435	27.21.aa.cdn.bp.final	100	100	Retrieved	07/28/10 at 08:47:22 PM EST	[remove]
409	57.21.nt123.ml.P2.final	250	250	Retrieved	07/28/10 at 04:14:57 AM EST	[remove]
408	57.21.dgn.ml.SupF.final	210	250	Running	07/28/10 at 04:12:37 AM EST	[remove]
407	57.21.dgn.ml.SubF_nLeuz.final	250	250	Retrieved	07/28/10 at 04:00:00 AM EST	[remove]
406	57.21.dgn.ml.SubF.final	250	250	Retrieved	07/28/10 at 03:58:29 AM EST	[remove]
405	57.21.dgn.ml.P.final	250	250	Retrieved	07/28/10 at 03:56:47 AM EST	[remove]
404	57.21.dgn.ml.F.final	250	250	Retrieved	07/28/10 at 03:49:17 AM EST	[remove]
403	57.21.dgn.ml.bgr.final	250	250	Retrieved	07/28/10 at 03:47:05 AM EST	[remove]
402	57.21.dgn.ml.bg.final	250	250	Retrieved	07/28/10 at 03:40:47 AM EST	[remove]
401	57.21.nt123.ml.supF.final	250	250	Retrieved	07/27/10 at 02:22:37 PM EST	[remove]
400	57.21.nt123.ml.supf_nLeuz.final	250	250	Retrieved	07/27/10 at 02:21:28 PM EST	[remove]
399	57.21.nt123.ml.subf.final	250	250	Retrieved	07/27/10 at 02:20:10 PM EST	[remove]
398	57.21.nt123.ml.P.final	250	250	Retrieved	07/27/10 at 02:18:21 PM EST	[remove]
397	57.21.nt123.ml.F.final	250	250	Retrieved	07/27/10 at 02:12:09 PM EST	[remove]

akawahar

[My account](#)[Log out](#)[Home](#) » [GARLI Web Service](#) – view job

GARLI Web Service – view job

[< back to job select](#)

Job Status

Job name: 27.21.aa.cdn.bp.final

Progress: 100/100

[\[Download zipped job files\]](#)

Job Files

- [27.21.14793.070910.nex](#)
- [garli.conf](#)

postprocessing

- [all.boot.tre](#)
- [tree_files](#)

Job 1

- [stdout](#)

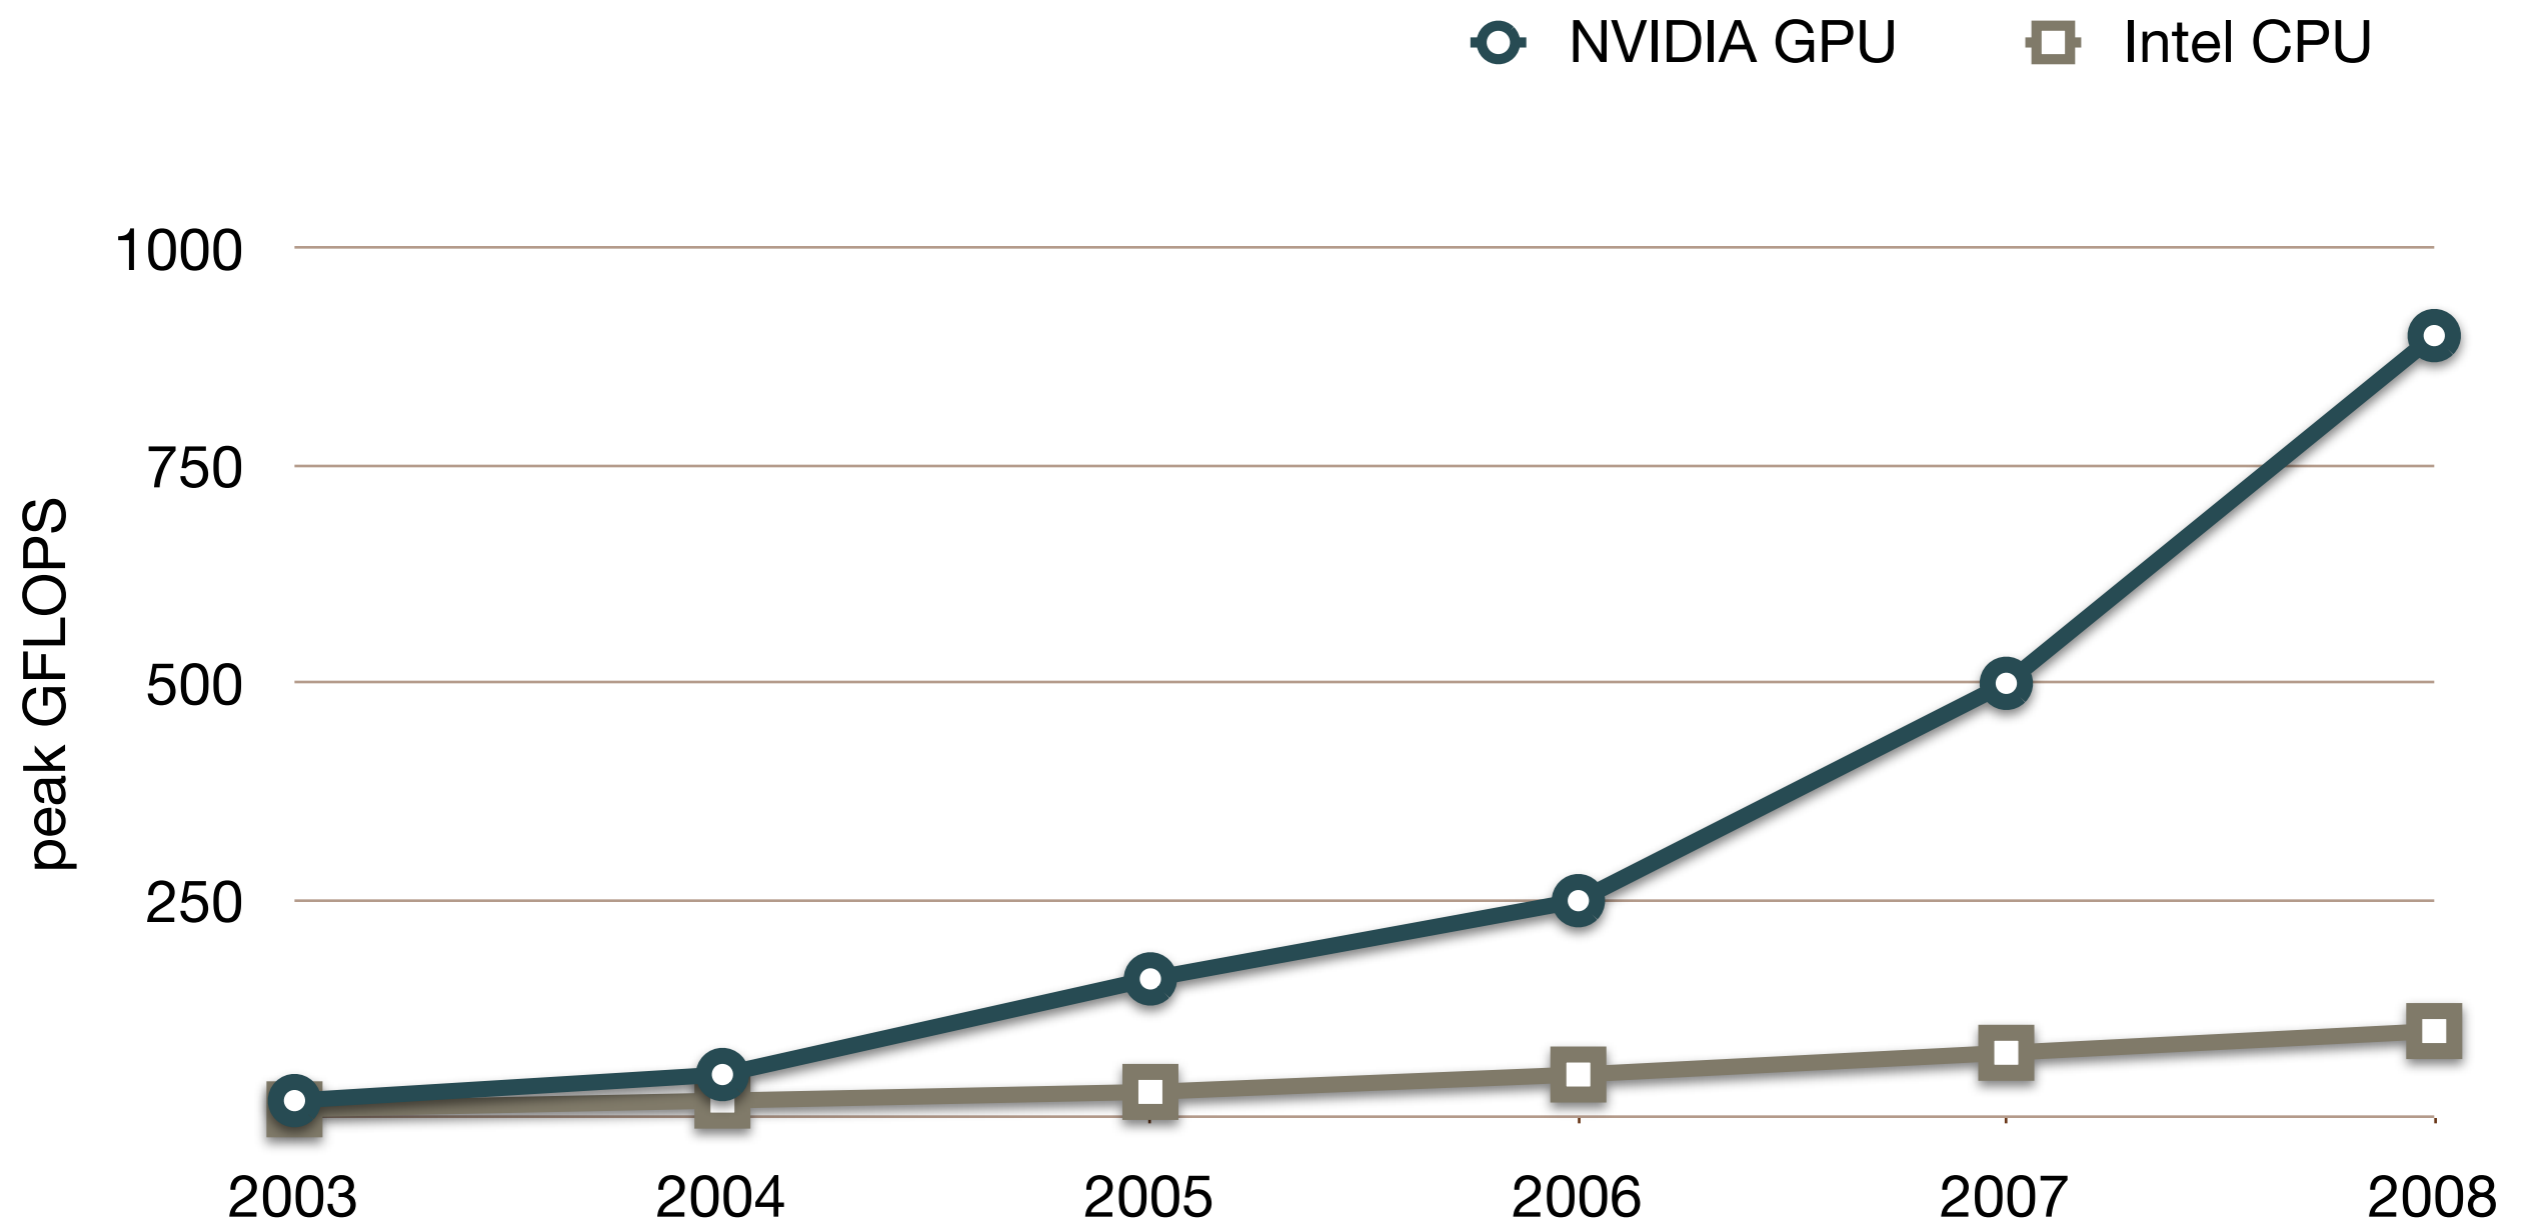
GPUs: background



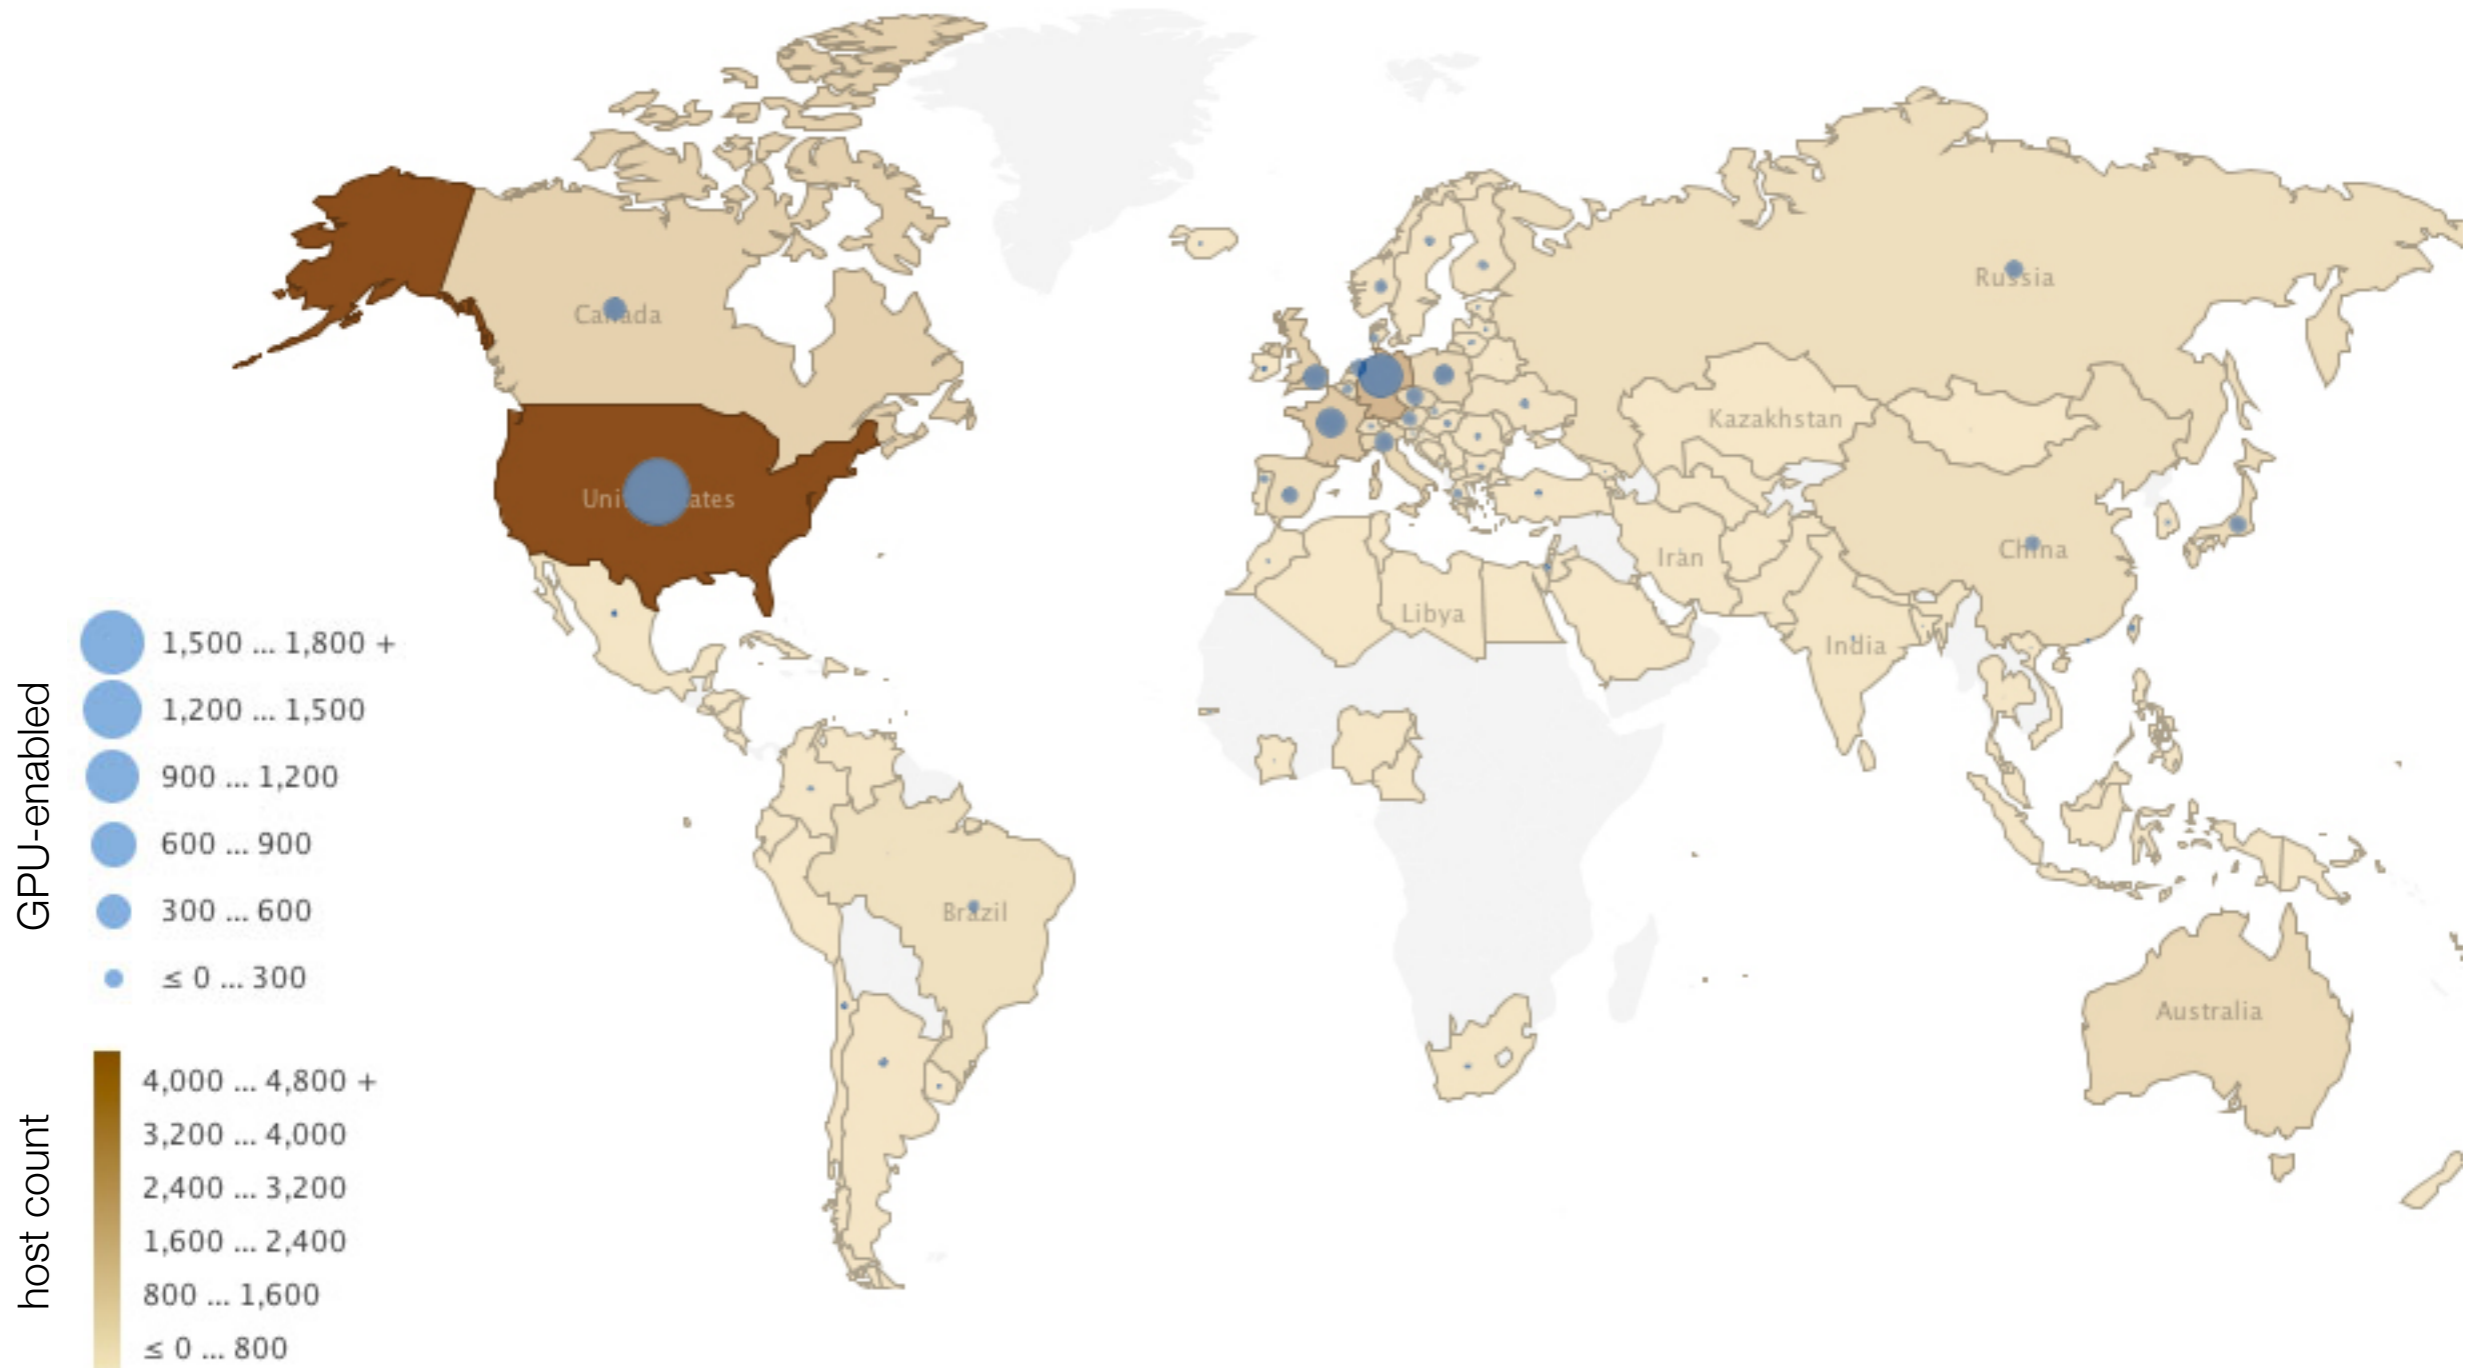
- GPUs are inexpensive, dedicated numerical processors originally designed for rendering 3D computer graphics
- GPUs contain 100s of processing cores on a single chip
- Each core carries out the same operations in parallel on different input data – single program, multiple data (SPMD) paradigm
- Extremely high arithmetic intensity if one can transfer the data onto and results off of the processors efficiently

GPUs: background

- general-purpose computing on GPUs




The Lattice BOINC Project: hosts by country




period of January 1st to August 23rd, 2010; total of 15081 hosts (5750 GPU-enabled); created on Many Eyes © IBM

The Lattice BOINC Project: GPU models



period of January 1st to August 23rd, 2010; total of 5750 GPUs; created on Many Eyes © IBM

GARLI-GPU: speedups



GARLI on BOINC: challenges and future work

- Checkpointing
 - GARLI only checkpoints in certain phases
- Progress bar updating
 - due to the stochastic GA and flexible stopping conditions, it can be difficult to ascertain progress
- Runtime estimates
 - providing accurate runtime estimates *a priori* is important for a variety of reasons; has been difficult to do so because of the high variance in types of workunits
solution: machine learning with random forest to estimate runtimes in advance
- Long jobs, short pieces
 - if we could break up workunits into shorter, even-length pieces, we could give our volunteers a more consistent experience, and would probably increase throughput

References

- Bazinet, A. L. 2009. **The Lattice Project: A Multi-model Grid Computing System.** Master's thesis, University of Maryland.
- Bazinet, A. L., and M. P. Cummings. 2008. **The Lattice Project: a Grid research and production environment combining multiple Grid computing models.** Pages 2-13. In Weber, M. H. W. (Ed.) *Distributed & Grid Computing - Science Made Transparent for Everyone. Principles, Applications and Supporting Communities.* Rechenkraft.net, Marburg.
- Myers, D. S., A. L. Bazinet and M. P. Cummings. 2008. **Expanding the reach of Grid computing: combining Globus- and BOINC-based systems.** Pages 71-85. In Talbi, E.-G. and A. Zomaya (Eds.) *Grids for Bioinformatics and Computational Biology, Wiley Book Series on Parallel and Distributed Computing.* John Wiley & Sons, New York.
- Bazinet, A. L., D. S. Myers, J. Fuetsch and M. P. Cummings. 2007. **Grid Services Base Library: a high-level, procedural application program interface for writing Globus-based Grid services.** *Future Generation Computer Systems* 22:517-522.
- Myers, D. S., and M. P. Cummings. 2003. **Necessity is the mother of invention: a simple Grid computing system using commodity tools.** *Journal of Parallel and Distributed Computing.* 63:578-589.
- Zwickl, D. J. 2006. **Genetic algorithm approaches for the phylogenetic analysis of large biological sequence datasets under the maximum likelihood criterion.** Ph.D. dissertation, The University of Texas at Austin.

More information



-  **THE LATTICE PROJECT** <http://lattice.umiacs.umd.edu/>
<http://boinc.umiacs.umd.edu/>
- Contacts:
 - Adam Bazinet (pknut777@umiacs.umd.edu)
 - Michael Cummings (mike@umiacs.umd.edu)