

# Some thoughts on HEP Computing in the era of Grids and Clouds

Miron Livny  
Computer Sciences Department  
University of Wisconsin-Madison  
[miron@cs.wisc.edu](mailto:miron@cs.wisc.edu)



As the Information Technology (IT) 'landscape' continues to change in an ever faster pace the HEP computing model has to continue to adapt so that it delivers in the most cost effective way the ever more demanding computing needs of the community.

The main IT forces that are changing this landscape today are:

- Multi-core and Graphics Processing Units (GPUs)
- Open source software (Hadoop, RH-MRG)
- Cloud Computing (Elastic Compute Cloud (EC2))
- Complexity and cost of software
- Broad adoption of distributed computing (grids)

# Future of Grid Computing

**CHEP 07**

Michael Lipp  
Computer Science Department  
University of Wisconsin-Madison  
miron@cs.wisc.edu



The Tulmod says in the name of Rabbi Yochanan,

"Since the destruction of the Temple, prophecy has been taken from prophets and given to fools and children."

(Baba Batra 12b)



**CHEP 07** [www.cs.wisc.edu/~miron](http://www.cs.wisc.edu/~miron)



# The Grid Computing Movement

I believe that as a movement grid computing ran its course.

- No more an easy source of funding
- No more an easy way to get the "troops" mobilized
- No more an easy sell of software tools
- No more an easy way to get your papers published or your press releases posted



**CHEP 07** [www.cs.wisc.edu/~miron](http://www.cs.wisc.edu/~miron)



# Distributed Computing

Distributed computing is here to stay and to continue to evolve as processing, storage and communication resources are more widely adopted and get more powerful and cheaper

- Big science is inherently distributed
- Most scientific disciplines (and many commercial sectors) depend on High Throughput Computing (HTC) capabilities



**CHEP 07** [www.cs.wisc.edu/~miron](http://www.cs.wisc.edu/~miron)



# Keynote 3: When All Computing Becomes Grid Computing

**Speaker:** Prof. Daniel A. Reed  
Chancellor's Eminent Professor  
Director, Renaissance Computing Institute  
University of North Carolina at Chapel Hill

**Now @ Microsoft**

## Abstract:

Scientific computing is moving rapidly from a world of "reliable, secure parallel systems" to a world of distributed software, virtual organizations and high-performance, though unreliable parallel and distributed systems with few guarantees of availability and quality of service. In addition, a tsunami of new experimental and computational data poses equally vexing problems in analysis, transport, visualization and collaboration. This transformation poses daunting scaling and reliability challenges and necessitates new approaches to collaboration, software development, performance measurement, system reliability and coordination. This talk describes Renaissance approaches to solving some of today's most challenging scientific and societal problems using Grids and parallel systems, supported by rich tools for performance analysis, reliability assessment and workflow management.

**CHEP 07**

One of the main accomplishments of the 'grid movement' is a close and productive partnership between Computer Scientists and the HEP community.

- One form of 'sharing' that underpins distributed computing
- What will it take to sustain and expand this partnership to provide the most cost effective computing capabilities to the HEP community and advance our understanding of distributed computing?

As a computer scientist (not a computing professional) I believe that we must understand the evolution of the HEP computing model (how we got to where we are today) as we develop short and long term plans for HEP computing. Unfortunately, very little attention (if any) has been devoted to establishing such an understanding.

- Few Computer Scientists at HEP laboratories
- No 'HEP-informatics'
- Focus on technology and 'fashionable trends' not principals

# My Background

- 1975 - B.Sc in Mathematics and Physics
- 1983 - PhD in Computer Science (Distributed systems)
- 1986 - First production deployment of the Condor distributed resource management system
- 1990 - Collaboration with the SMC group at NIKHEF
- 1992 - Seminar at CERN on a world wide distributed computing environment

# Global Scientific Computing via a Flock of Condors

# CERN 92

*Miron Livny*

Computer Sciences Department  
University of Wisconsin — Madison  
Madison, Wisconsin  
{miron@cs.wisc.edu}

## MISSION

Give scientists effective and efficient access to large amounts of cheap (if possible free) CPU cycles and main memory storage

## THE CHALLENGE

How to turn existing privately owned clusters of *workstations, farms, multiprocessors, and supercomputers* into an efficient and effective Global Computing Environment?

In other words, how to minimize wait while idle?

## APPROACH

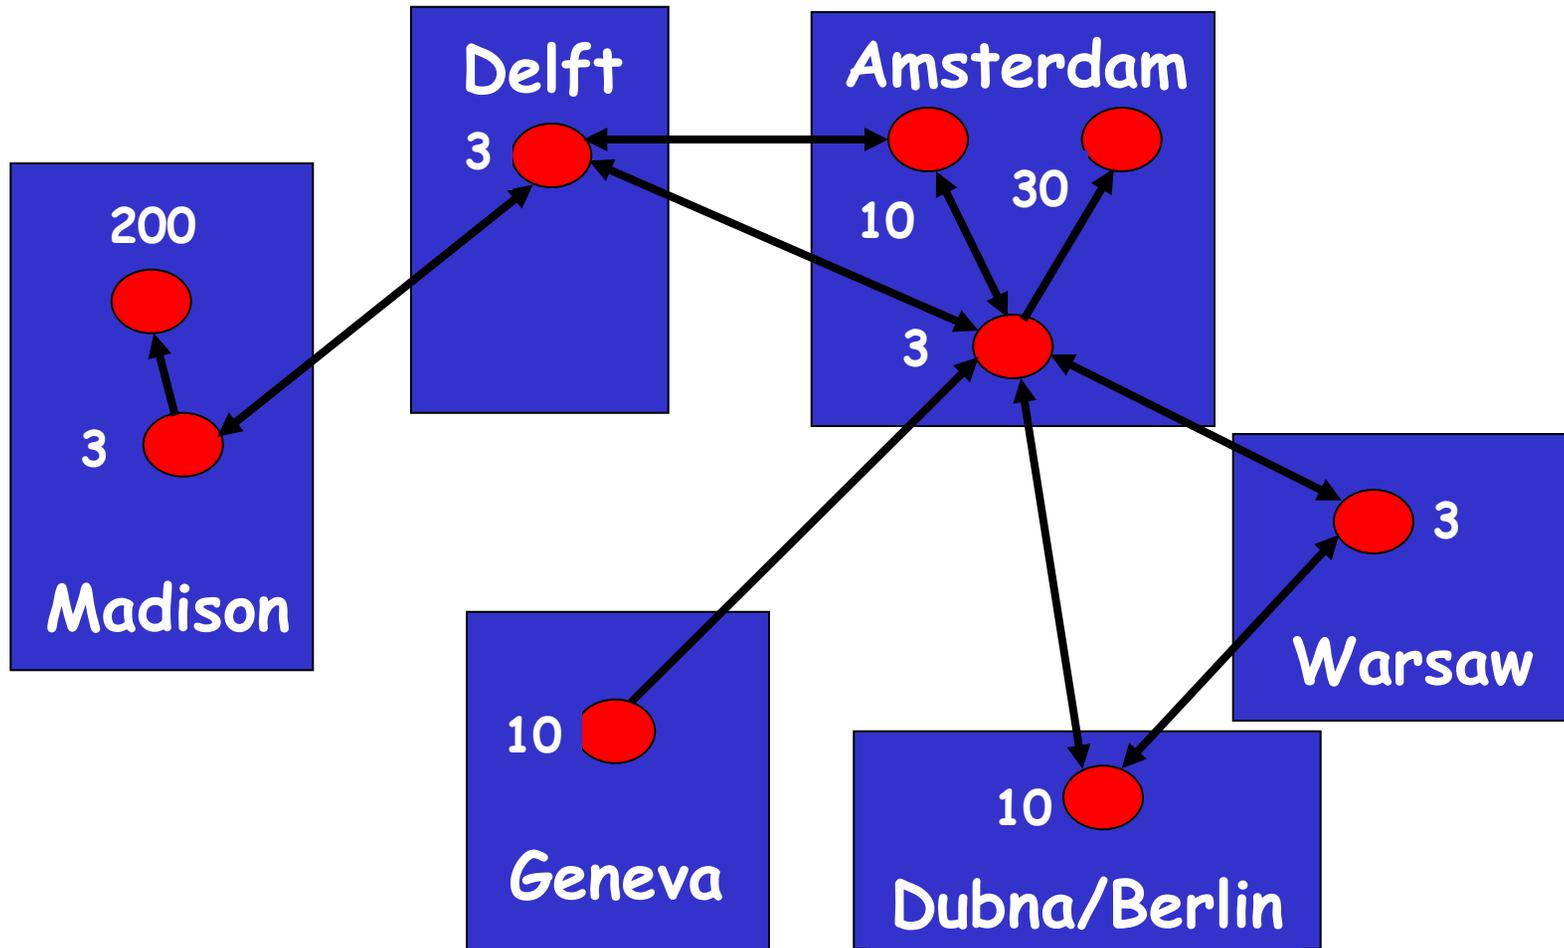
Use wide-area networks to transfer batch jobs between Condor systems

- Boundaries of each Condor system will be determined by physical or administrative considerations

## TWO EFFORTS

- UW CAMPUS**  
Condor systems at Engineering, Statistics, and Computer Sciences
- INTERNATIONAL**  
We have started a collaboration between CERN-SMC-NIKHEF-Univ. of Amsterdam, and University of Wisconsin-Madison

# The 1994 Worldwide Condor Flock



# Background

- 1996 - Condor deployed at INFN and collaboration with WA92 group in Bologna
- 1999 - Joined the Particle Physics Data Grid (PPDG) project (Computer Science PI on 2002)
- 2000 - Joined the Grid Physics Network (GriPhyN) project
- 2000 - Collaboration with UW-CMS group (today a CMS Tier-2)
- 2001 - Joined the international Virtual Data Grid laboratory (iVDGL) project
- 2003 - Collaboration to build the CDF CAF on Condor
- 2005 - Collaboration with UW-ATLAS (today a Tier-3)
- 2006 - PI of the Open Science Grid (OSG)
- 2007 - Condor distributed as part of RHEL



# Why do I collaborate with HEP?

- > The HEP community is committed to the distributed computing model
- > Provides a one-of-a-kind worldwide distributed computing 'laboratory'
- > A demanding, engaged and organized user community

# Today I am ...

- Technology Director of the Open Science Grid (OSG)
- Condor is distributed as part of Fedora and RHEL and supported by RedHat
- Both US Tier-1s and many Tier-2s and Tier-3s are using Condor
- Our technologies are part of the software stack of CDF, DZero, ATLAS and CMS as well as EGEE and OSG
- My group is responsible for the Virtual Data Toolkit (VDT)
- We collaborate with RH and Amazon on developing capabilities for the Elastic Computing Cloud (EC2)





Conference started!

Days and 12 hours

## List of Panels

to

Dates

is

on

Committee

on committee

workshop

- Declarative, Domain-Specific Languages - Elegant Simplicity or a Hammer in S of a Nail?  
Sam Madden, *MIT*; Alan Demers, *Cornell University*; Michael Carey, *BEA*; Boon Loo, *University of Pennsylvania*; John Whaley, *moka5*
- Scientific Data Management: An Orphan in the Database Community?  
Randal Burns (moderator), *Johns Hopkins University*; Susan B. Davidson, *Cornell University*; Yannis Ioannidis, *University of Athens*; Miron Livny, *University of Wisconsin-Madison*; Jignesh M. Patel, *University of Michigan*
- Cloud Computing-Was Thomas Watson Right After All?  
Raghu Ramakrishnan (moderator), *Yahoo!*; Eric Baldeschwieler, *Yahoo!*; James Hamilton, *Microsoft*; Miron Livny, *University Wisconsin-Madison*; Yossi Matias, *Google*; Hamid Pirahesh, *IBM*





# Condor in the Clouds



Conference started!

Days and 12 hours

## List of Panels

# Clouds

to

Dates

is

on

Committee

on Committee

website

- **Declarative, Domain-Specific Languages - Elegant Simplicity or a Hammer in S of a Nail?**  
Sam Madden, *MIT*; Alan Demers, *Cornell University*; Michael Carey, *BEA*; Boon Loo, *University of Pennsylvania*; John Whaley, *moka5*
- **Scientific Data Management: An Orphan in the Database Community?**  
Randal Burns (moderator), *Johns Hopkins University*; Susan B. Davidson, *Cornell University*; Yannis Ioannidis, *University of Athens*; Miron Livny, *University of Wisconsin-Madison*; Jignesh M. Patel, *University of Michigan*
- **Cloud Computing-Was Thomas Watson Right After All?**  
Raghu Ramakrishnan (moderator), *Yahoo!*; Eric Baldeschwieler, *Yahoo!*; James Hamilton, *Microsoft*; Miron Livny, *University Wisconsin-Madison*; Yossi Matias, *Google*; Hamid Pirahesh, *IBM*



The words of Koheleth son of David, king in Jerusalem ....

Only that shall happen  
Which has happened,  
Only that occur



Conference started!

List of Panels

Days and 12 hours

Which has occurred;  
There is nothing new  
Beneath the sun!

to

Dates

is

on

Committee

on Committee

Ecclesiastes Chapter 1 verse 9

Overview

- Declarative, Domain-Specific Languages - Elegant Simplicity or a Hammer in Search of a Nail?  
Sam Madden, MIT; Alan Demers, Cornell University; Michael Carey, BEA; Boon Loo, University of Pennsylvania; John Whaley, moka5
- Scientific Data Management: An Orphan in the Database Community?  
Randal Burns (moderator), Johns Hopkins University; Susan B. Davidson, Cornell University; Yannis Ioannidis, University of Athens; Miron Livny, University of Wisconsin-Madison; Jignesh M. Patel, University of Michigan
- Cloud Computing-Was Thomas Watson Right After All?  
Raghu Ramakrishnan (moderator), Yahoo!; Eric Baldeschwieler, Yahoo!; James Hamilton, Microsoft; Miron Livny, University Wisconsin-Madison; Yossi Matias, Google; Hamid Pirahesh, IBM





# Our view of a cloud

An autonomous computing (processing, storage and networking) resources with an interface that supports remote invocation of "jobs" and staging of input/output data.

- Looks and feels like any other grid site
- Likely to have proprietary APIs
- Likely to have different cost models
- Likely to have different SLAs
- Likely to have different usage policies

*Declarative, Domain-Specific Languages: Elegant Simplicity or a Hammer in Search of a Nail?*  
Sam Madden, MIT; Atan Demers, Cornell University; Michael Carey, BEA; Boon Boon, University of Pennsylvania; John Whaley, moka5

*Scientific Data Management: An Orphan in the Database Community?*  
Kunal Bhandal, Johns Hopkins University; John B. Davidson, Cornell University; Yannis Ioannidis, University of Athens; Miron Livny, University of Wisconsin-Madison; M. Paell, University of Michigan

*Cloud Computing-Was Thomas Watson Right After All?*  
Rashu Rana, Finnegan (moderator), Yahoo!; Eric Baldeschwieler, Yahoo!; James Hamilton, Microsoft; Miron Livny, University Wisconsin-Madison; Yossi Matias, Google; Hamid Pirahesh, IBM



# What do we do with clouds?

- > Turn VMs into "first class citizens" in the Condor framework
- > Interact with users in academia and industry who express interest in using computing resources offered by clouds
- > Add EC2+S3 to the (long) list of remote resources Condor can harness (or delegate work to)
- > Explore possible enhancements to our matchmaking and workflow technologies to support provisioning of cloud resources (including inter-cloud migration)
- > Understand the semantics of the EC2+S3 services, protocols and infrastructure so that we can provide a Condor "overlay" that expend local capabilities to include these resources
- > Monitor new cloud "formations"

conference started!

rs and 12 hours

Date:

is

on

committee

on Committee

workshop

Declarative, Domain-Specific Languages - Elegant Simplicity or a Hammer in S of a Nail?

Sam Madden, MIT; Hans-Joerg Boehm, Cornell University; Michael Carey, Penn State; Loo, University of Pennsylvania; John Whaley, moka5

Scientific Data Management: An Update from the Database Community?

Ronald Kim (moderator), Johns Hopkins University; Susan B. Davidson, Cornell University; Yannis Ioannidis, University of Athens; Miron Livny, University of Wisconsin-Madison

Cloud Computing - Was Thomas Watson Right After All?

Raghu Ramakrishnan (moderator), Yahoo!; Eric Baldeschwieler, Yahoo!; James Hamilton, Microsoft; Miron Livny, University Wisconsin-Madison; Yossi Matias, Google; Hamid Pirahesh, IBM



# Cloud Computing

- > Part of the mix of computing resources and services that needs to be integrated into our computing framework
- > Defines a cost model for computing capabilities
- > Promotes Virtualization technologies
- > Encourages "on-the-fly" deployment of software stacks

How can we accommodate  
an unbounded  
need for computing with  
an unbounded  
amount of resources?

“ ... Since the early days of mankind the primary motivation for the establishment of *communities* has been the idea that by being part of an organized group the capabilities of an individual are improved. The great progress in the area of inter-computer communication led to the development of means by which stand-alone processing sub-systems can be integrated into multi-computer *communities*. ... ”

Miron Livny, “ *Study of Load Balancing Algorithms for Decentralized Distributed Processing Systems.*”,  
Ph.D thesis, July 1983.

