

Introduction to Condor

Alain Roy

Associate Researcher Condor Project

What is Condor?

University Wisconsin-Madison





~20 full-time staff

~12 students

Production-quality software

Research

Condor software

Other software

This talk: Just
Condor software

High-throughput batch system



Batch System:
Run program on
many computers

High-Throughput:

FLOPS / Year

Not

FLOPS/Second

Emphasis on
reliability

No emphasis on
overheads for
short-running jobs

Condor does not
optimize job startup
(10s of seconds)

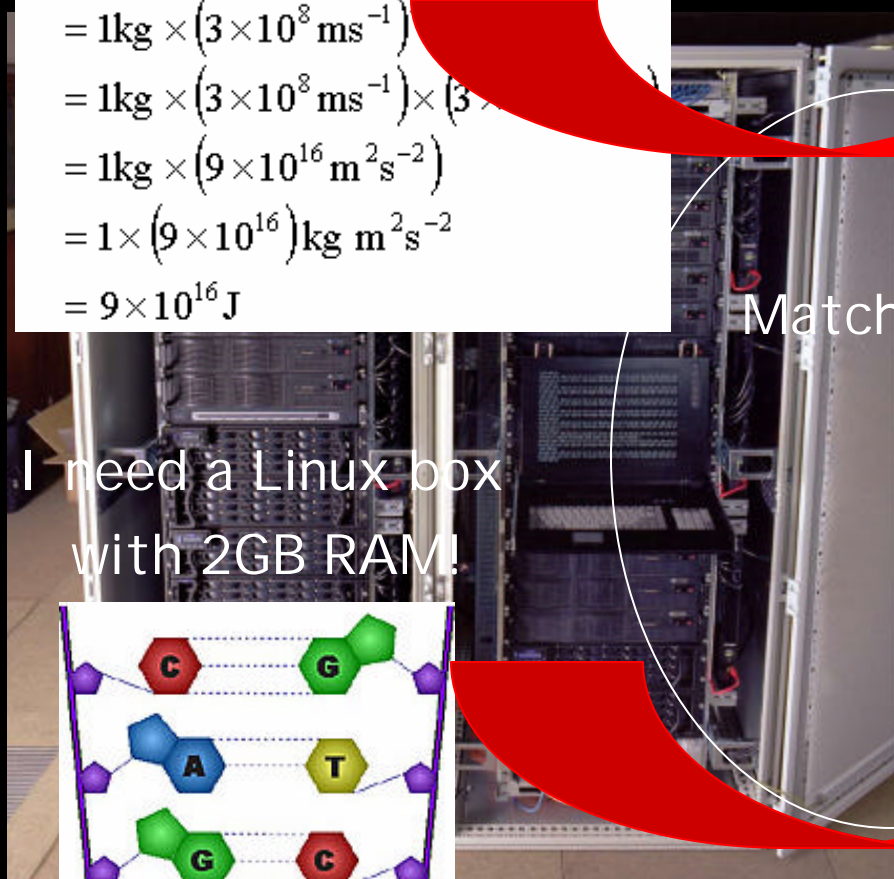
Long-term
performance

Condor Takes Computers...


I need a Mac!

$$E = mc^2$$
$$= 1\text{kg} \times (3 \times 10^8 \text{ms}^{-1})^2$$
$$= 1\text{kg} \times (3 \times 10^8 \text{ms}^{-1}) \times (3 \times 10^8 \text{ms}^{-1})$$
$$= 1\text{kg} \times (9 \times 10^{16} \text{m}^2 \text{s}^{-2})$$
$$= 1 \times (9 \times 10^{16}) \text{kg m}^2 \text{s}^{-2}$$
$$= 9 \times 10^{16} \text{J}$$

Desktop Computers



I need a Linux box with 2GB RAM!



Matchmaker



Quick Definitions

Cluster: A dedicated set of computers not for interactive use



Pool:

A collection of computers
used by Condor

Dedicated or interactive

Matchmaking

Core Feature in Condor

Matchmaking
is Two-Way

A job has
requirements

“I need Linux and
2GB of RAM”

A computer has
requirements

“I only run jobs at
night”

Matchmaking
allows preferences

I need Linux

I prefer to have lots
of memory

Why two-way
matching?

Condor
conceptually has
three kinds of users

- Job Submitters
- Machine Owners
- Pool (Cluster) administrators

These may overlap... or they might not

Machine Owners
have preferences
like:

Prefer users from a
particular group

Prefer one job type
over another

System admins
have different
preferences

When can jobs
preempt other
jobs?

Which users are
high priority?

ClassAds



Fundamental to
matchmaking


ClassAds state
facts

“My job’s
executable is
analysis.exe”

“My machine’s load
is 2.1”

ClassAds state
preferences

“I require a Linux
computer”



ClassAds are:

Semi-structured

User-extensible

Schema-Free

Example:

```
MyType           = "Job"
TargetType       = "Machine"
ClusterId        = 1377
Owner            = "roy"
Cmd              = "analysis.exe"
Requirements     =
    (Arch == "INTEL" )
&& (OpSys == "LINUX" )
&& (Disk >= DiskUsage)
&& ((Memory * 1024) >= ImageSize)
...
```

Condor imposes
some schema

“Owner is a string”

Users can extend
the schema

```
AnalysisJobType = "simulation"
```

```
HasJava_1_4 = TRUE
```

ShoeSize = 11

Matchmaking can
use all attributes

```
Requirements =  
OpSys = "Linux"  
      &&  
HasJava_1_4 = TRUE
```

Submitting Jobs

Job is described
with a ClassAd

A queue per
submission
computer

Queues are *not*
centralized


Machine is
described with a
ClassAd

Matchmaker
compares
ClassAds

Condor & Grids



Submit to Globus & Other systems



Condor-G

Why?

Condor-G provides
reliability

Persistent queue,
fault tolerance...

Condor-G provides
flexibility

Policies to control
job execution
Methods to run
many jobs...

Who uses Condor?

Nanohub uses Condor-G
for grid jobs

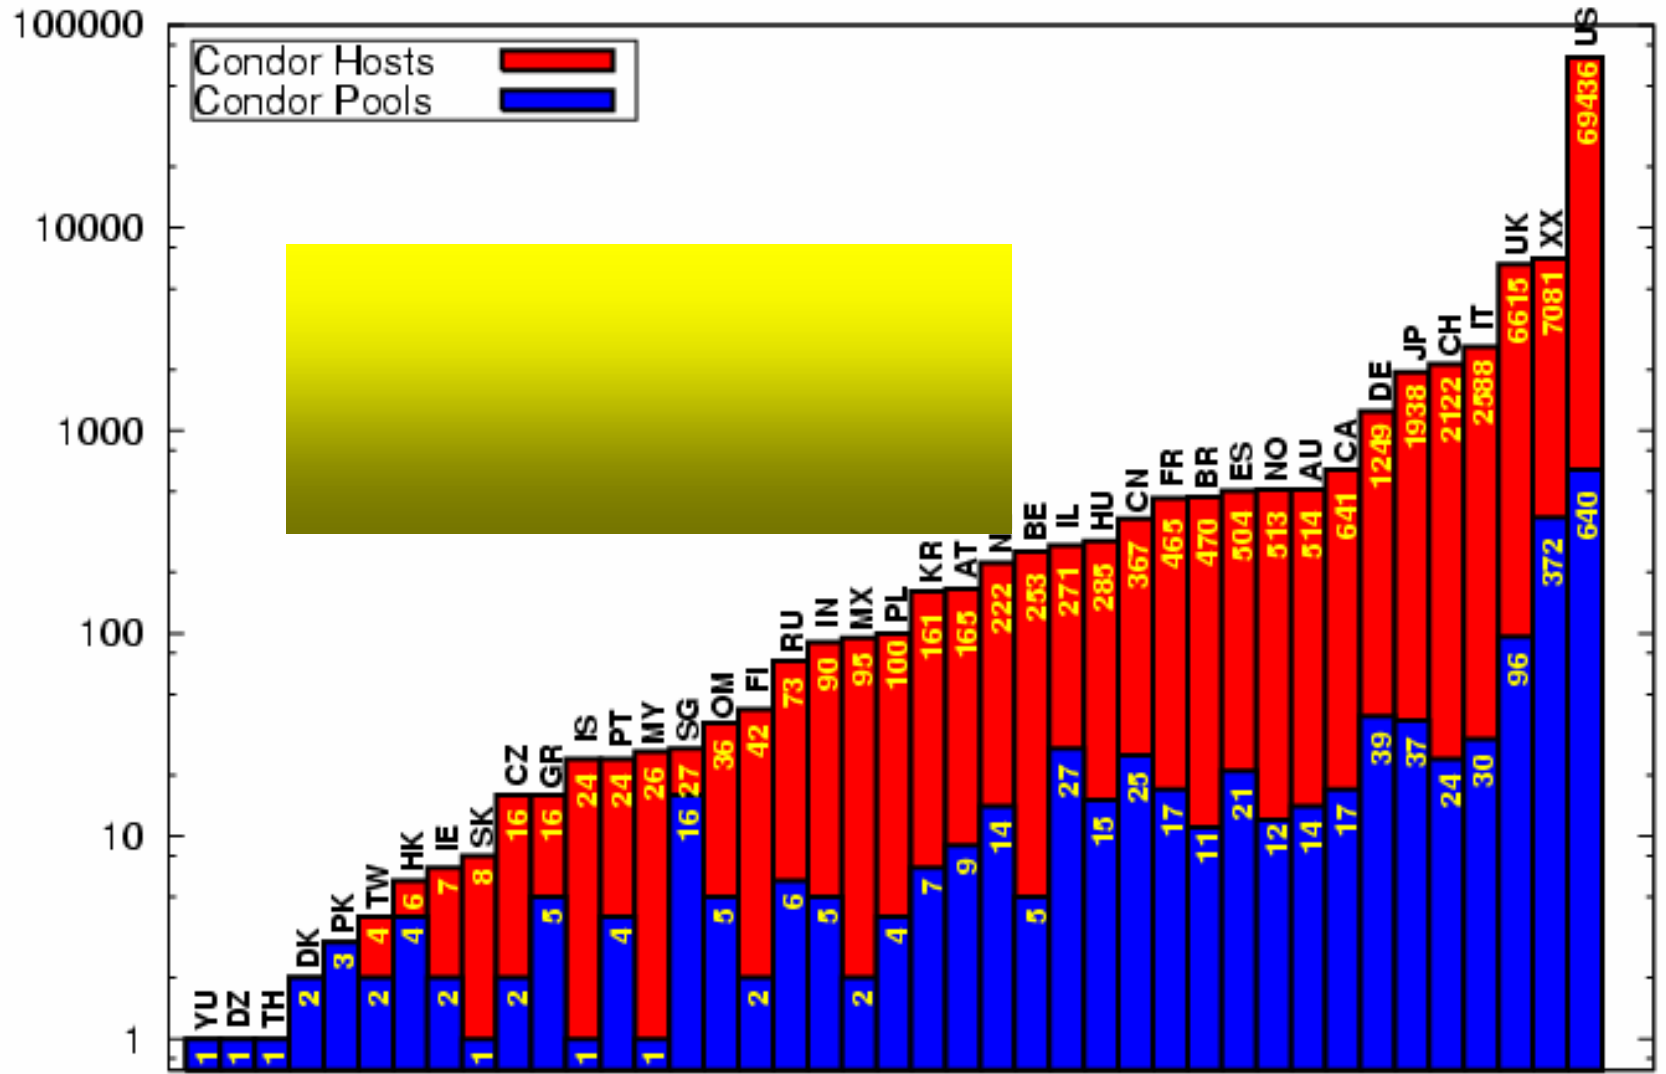
Research


National Laboratories

Companies

Known Condor Pools and Hosts by Country

Tue Jan 31 18:44:13 CST 2006





System Type:

- A AIX
- D DUX
- M OSX
- S SUNOS
- H HPUX
- I IRIX
- L Linux
- W Windows

CPU Type:

- I Intel x86
- 6 Intel IA64
- A Alpha
- P Power PC
- S SPARC
- H HPPA
- M MIPS


Size of Pool:

- > 1000 hosts
- 100-999 hosts
- 10-99 hosts
- 1-9 hosts

Legend:

- Software Version 6.V**: S, C, V, ∅
- Data Source**: O email, / udp

Learn more



www.cs.wisc.edu/condor

Software downloads



Manuals



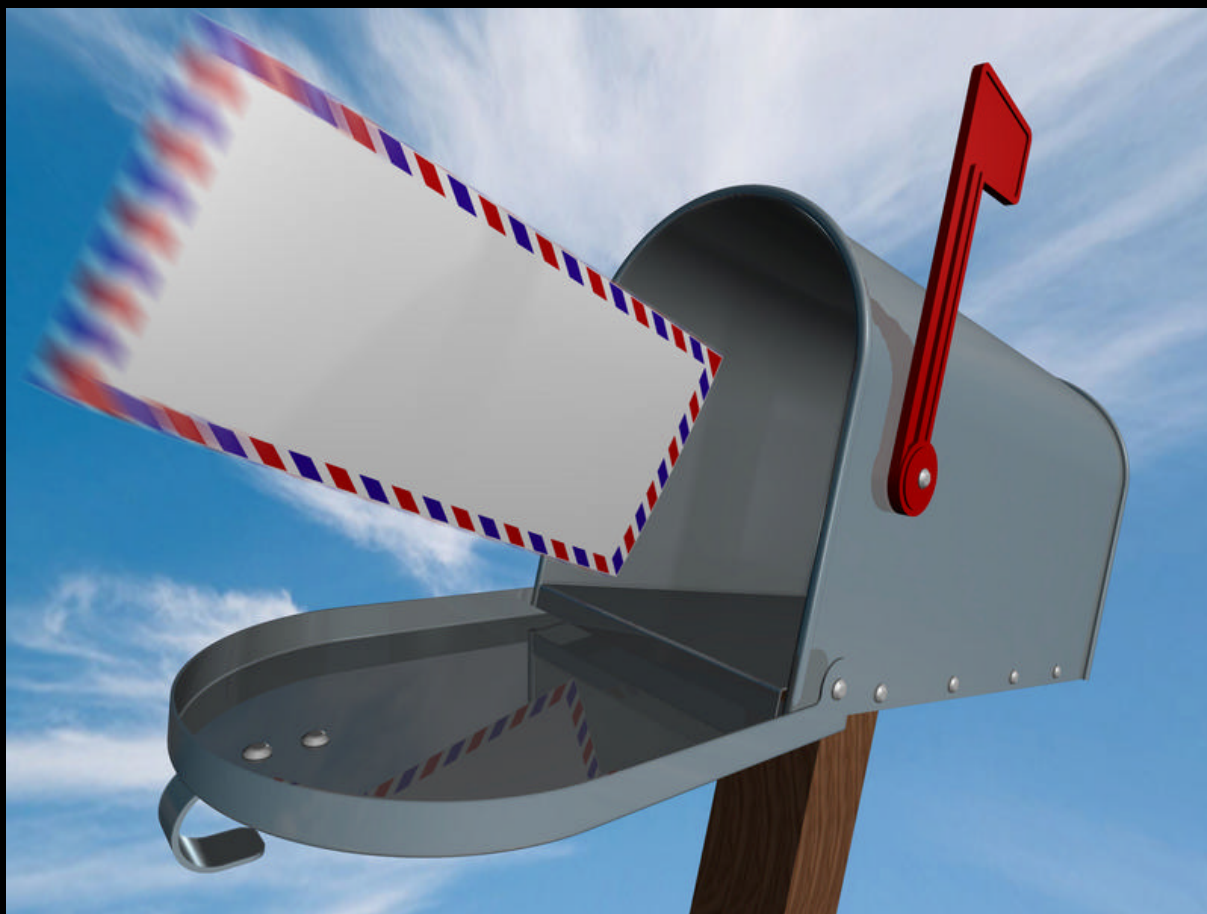
Research papers



Tutorials



Mailing Lists



www.cs.wisc.edu/condor