



Java is a trademark of Sun Microsystems, Inc.



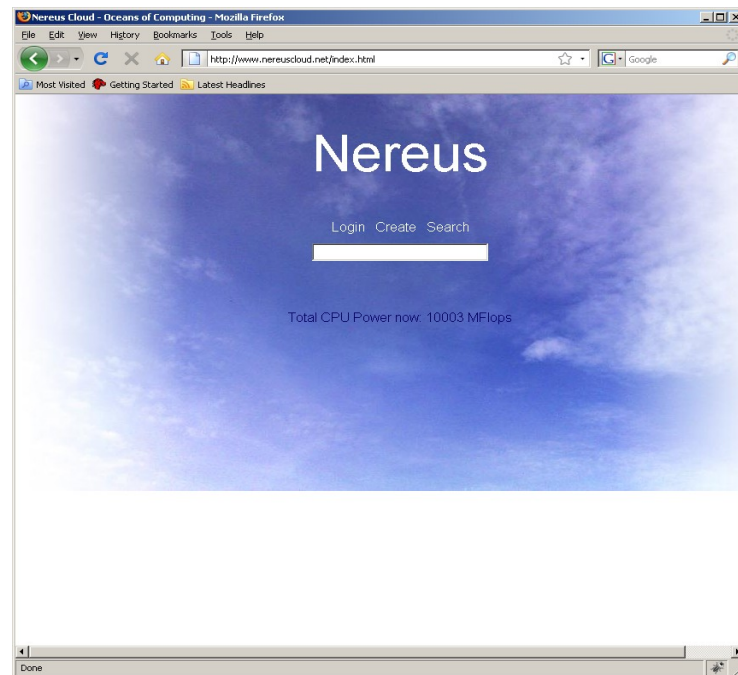
JavaOneSM

Nereus-V: Massively Parallel Computing of, by, and for the Community

Rhys Newman & Ian Preston
Oxford University
TS-5136

NereusV – Massively Parallel Computing



- > **NereusV** (Nereus 5) - 100% Java Desktop Grid computing technology
- > Donate idle CPU time by just clicking a link
- > Runs completely inside the applet sandbox
- > Firewall friendly



Prototype web portal at <http://www.nereuscloud.net>

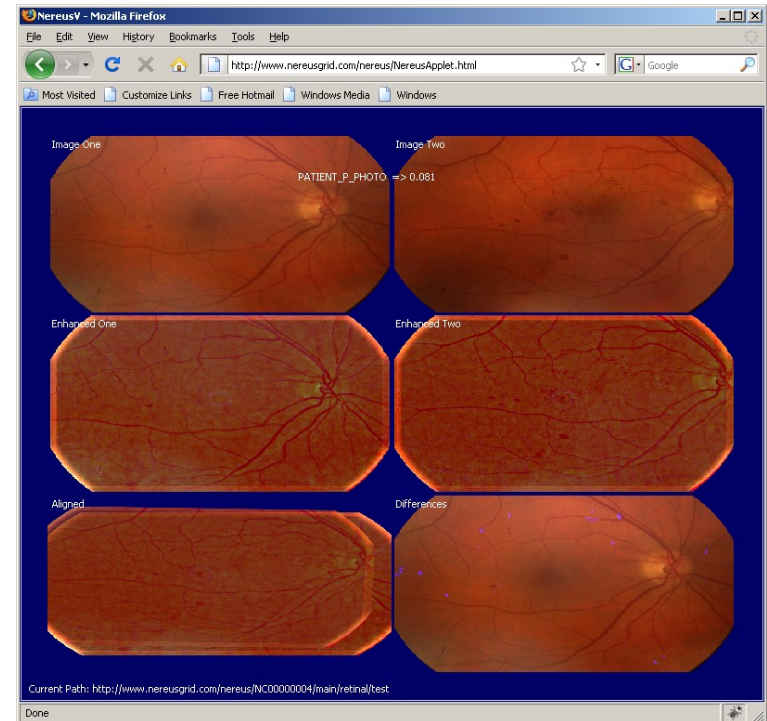
Demonstration of NereusV: Medical Image Processing

- > Changes in the retina indicates progress of eye disease
 - Diabetes sufferers particularly susceptible
- > Take images every few months and look for changes, but...
 - The retina will be photographed in different positions, by different cameras and operators
 - Some changes are just dust on the camera lens
 - Examining images is skilled work → \$\$\$
- > Use computing to “brute force” through pixels and detect changes.
 - 1000s CPU hours needed



Live Nereus Demo #1

1. Point browsers at www.nereusgrid.com
2. Use NereusCloud to push the image processing code out to donors
3. Collect the results!




Feel free to try it now with your laptop...

A Global Desktop Grid?

1Billion desktop PCs = 100 x (Top 500 supercomputers combined)

“Vanilla” computing power ever more in demand...



Analysis from Top500.org (Nov 2008)

- > BOINC is largest example with 550,000 computers
 - Runs over 30 world-leading (famous) research projects
 - Releases more computational power than the worlds fastest supercomputer.
- > But...
 - has only captured 0.05% of the global potential!
 - Barriers are still too high...

Why we've been waiting for NereusV

Security (Donors)

- PC owners are rightly scared of executing unvetted code

Convenience (Donors / Administrators)


- Installing software can be difficult
- Hardware / OS variations make support difficult
- Extra software → extra maintenance

Usability (Users)

- Develop/support code for a wide range of HW/SW

Scalability (Users / Administrators)

- Can the system grow to millions of clients?



Java Technology inherently solves most of these – NereusV solves the rest...




NereusV



- > Developed for over 4 years at Oxford University
 - Open Source – GPLv2
 - Pure Java – runs in any 1.6 VM
- > Easy!
 - Donate – one click does the trick
 - Virtually zero administration/maintenance
 - Easy to develop – as simple as Java Applets
- > Designed for global scale

NereusV – The Community Model



NereusV Service

In Package org.nereus.client ...

```
public interface NereusService
{
    public void init(ServiceContainer container);

    public void handleConnection(InetAddress remoteAddress,
                                HTTPRequest request,
                                HTTPResponse response) throws IOException;
}

public class ServiceContainer
{
    . . .
    public void accept(int timeout) throws IOException
    . . . .
}
```

There's AbstractService and AbstractListeningService to make life easier....

NereusV Server

> Proxies client connections to the internet

- The applet sandbox only lets the applets talk to the server from which they were downloaded
- The applet can only load classes with the same codebase, so the Nereus Server supports URL proxying:

`http://myserver/mycode.class` →

`http://nereusserver/urlproxy/myserver/mycode.class`

Note: numeric IP addresses not allowed, permits DNS management through `opendns`

- Network rules in a config file to prevent network abuse (NereusV firewall).

> Designed for 1000s of permanently connected clients

- Nereus 4 used NIO and limited threads
- NereusV now uses 1 thread per connection!


> Supports all services through HTTP

- Uses port 80 with different services distinguished by HTTP header content



NereusCloud – Top Level Service

- > Management services
 - Create an account which can be “pointed” to by Nereus Servers
 - Conveniently deploy/undeploy application software in one click
- > Does *not* do processing or network proxying
 - Very little load, so one logical site is sufficient for global scale



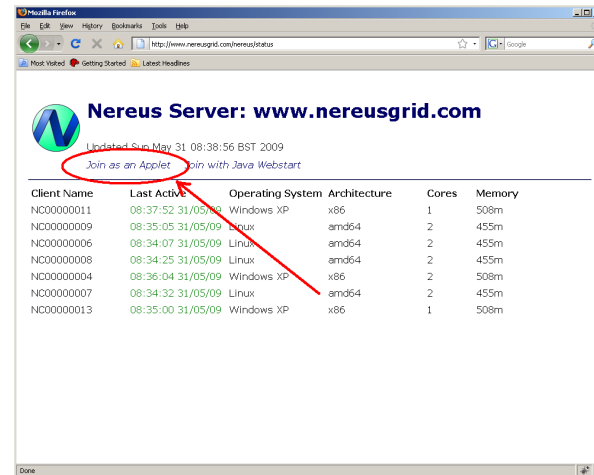
Live Demo #2 – Financial Modelling

- > Look for patterns in financial time series
 - Predicting outcomes better than 50/50 makes money
 - Hedge Funds pay a lot of smart people a lot of money to do this
- > Can we do better with dumber algorithms but with loads of computing power?

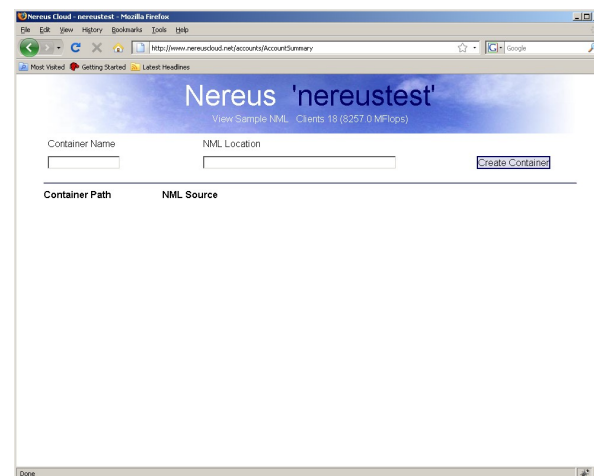


Live Demo #2 – Predicting the Dow Jones

- > Open a browser to www.nereusgrid.com
 - Click on “Join Server...”
- > We're in NereusCloud account *nereustest*
- > Lets turn to the live display to see the modelling in action....



Client Name	Last Active	Operating System	Architecture	Cores	Memory
NC00000011	08:37:52 31/05/09	Windows XP	x86	1	508m
NC00000009	08:35:05 31/05/09	Linux	amd64	2	455m
NC00000006	08:34:07 31/05/09	Linux	amd64	2	455m
NC00000008	08:34:25 31/05/09	Linux	amd64	2	455m
NC00000004	08:36:04 31/05/09	Windows XP	x86	2	508m
NC00000007	08:34:32 31/05/09	Linux	amd64	2	455m
NC00000013	08:35:00 31/05/09	Windows XP	x86	1	508m



Nereus 'nereustest'

View Sample NML Clients: 18 (8257.0 MFlops)


Container Name: NML Location:

Container Path: NML Source:


NereusV is *Green IT*

Nereus does not imply machines are on for longer – there are 2 types of *Idle Time*

- > Type I: No-one is sitting in front of the machine and its not running as a server



- > Type II: A user is using the machine but not demanding a lot of CPU time




- > Even if machines are turned off to minimise type I idle time, there's plenty of type II for Nereus!

NereusV is *Green IT*

Using Nereus *compensates* for the environmental impact of a desktop PC


- > PC Manufacture is 50% of its lifetime CO₂ footprint
- > Use the cloud?
 - Cooling is a 30-100% overhead
 - Datacentres have a large CO₂ footprint
- > Use NereusV?
 - Uses existing computers
 - No extra cooling or housing costs
 - Fewer servers in datacentres → less CO₂



SETI@Home, Folding@Home in Java!

> Oxford Physiologists using NereusV for Molecular modelling

- Atomic level simulation of protein agglomeration in Alzheimer's research
- Results submitted to Physical Review E



> Working with Oxford Astrophysics for sky surveys

- Not quite SETI@Home - looking for pulsars rather than ET!
- Like SETI@Home there's a *lot* of data which can be processed in parallel



NereusV + JPC – Running Native Code

- > NereusV only runs Java code...
...and then only within the Applet sandbox
- > What about legacy code, most often x86 Linux?
- > Use JPC: a pure Java x86 PC emulator....

JPC Launched - JavaOne 2007




JPC Improved - JavaOne 2008



JPC – Java x86 Hardware Emulator

CD-ROM Drive



Network

Interrupt Controller

PCI Host Bridge

PCI ISA Bridge

PS/2 Interface

DMA Controller

Interval Timer

PCI Bus

Processor

VGA Graphics

VGA BIOS

Real Time Clock

MMU




Keyboard

Mouse



Floppy Drive



BIOS

IDE Interface

Memory

Floppy Controller

Hard Disk Drive



JPC Demo #1

- > JPC Runs faster and smoother
 - Doom playable in an applet...




But there's more....



JPC - Recent Progress









- > JPC Runs much faster and more smoothly
 - No need for signed applet any more
 - Better dynamic compilation scheduling
- > Hardware level snapshots
 - Virtual PC is “Instant On” wherever downloaded
 - Instant hardware-level pause/resume
- > Using the latest Java 6 plugin we finally get more memory
 - Far more reliable startup and restart:
reboot = refresh!





JPC - Recent Progress


- > Loads of bug fixes emulation much more accurate
- > Currently testing with 23 Linux distributions including:

OS Name		Desktop	Prompt	Not Booting
DSL		✓	✓	
DELI			✓	
Feather		✓	✓	
Open BSD			✓	
TTY			✓	
Ubuntu 8.10				✗ Kernel Panic, SMP issue?
Puppy				✗ Kernel Panic, fatal exception in interrupt
Knoppix				✗ Kernel Panic, fatal exception in interrupt



Multi-GB Disk Images?


- > **Store disk images on a server**
 - Data loaded over the network on demand
 - Writes are cached locally
 - Better load time and memory footprint
- > **Use HTTP 1.1 range header to request parts of the Disk image**
 - Remote HD images can be served by *any* standard webserver e.g. Amazon....
- > **The guest OS is unaware of this, it thinks it's got a normal HDD!**





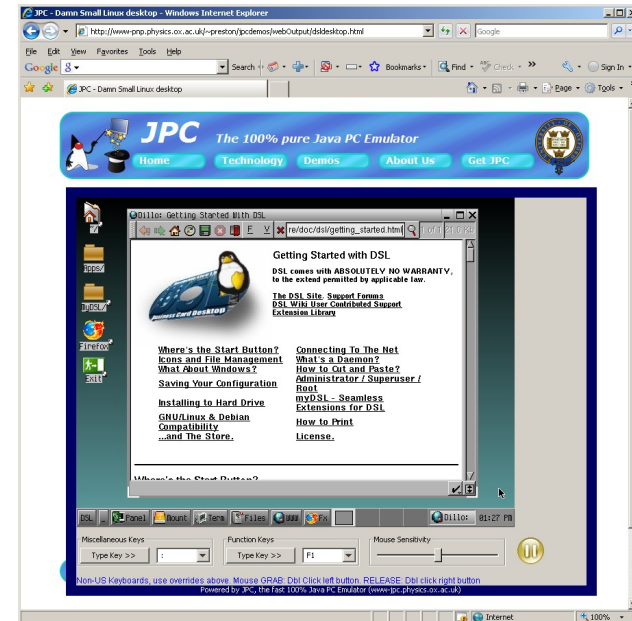
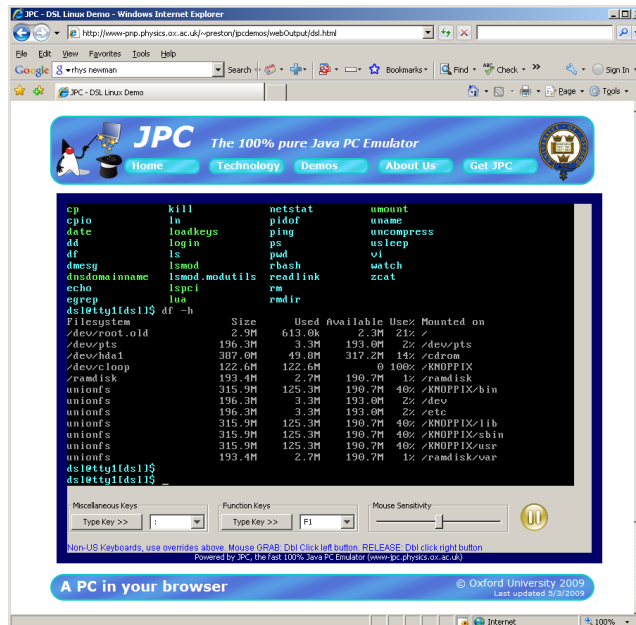
Task Specific Compilation

- > **Dynamic compiler builds classes at runtime to emulate blocks of x86 code**
 - Special classloader loads these into the JVM
➔ 10x speed increase at least
 - Custom classloader not allowed in the applet sandbox!
- > **Run a task/program/OS in the JPC Application...**
 - Save the custom classes to disk in a JAR file
 - Include this JAR in the classpath of any applet
 - Class names hashed from the x86 code
- > **Compiled classes even inside the sandbox**
 - JPC Applet even faster without the background compiler running...



JPC Demo #2


- “Instant On” Linux inside the Java Applet sandbox!
- Remote HDD
- Snapshot taken at the prompt/desktop



Visit <http://www-jpc.physics.ox.ac.uk/> or just Google “JPC”

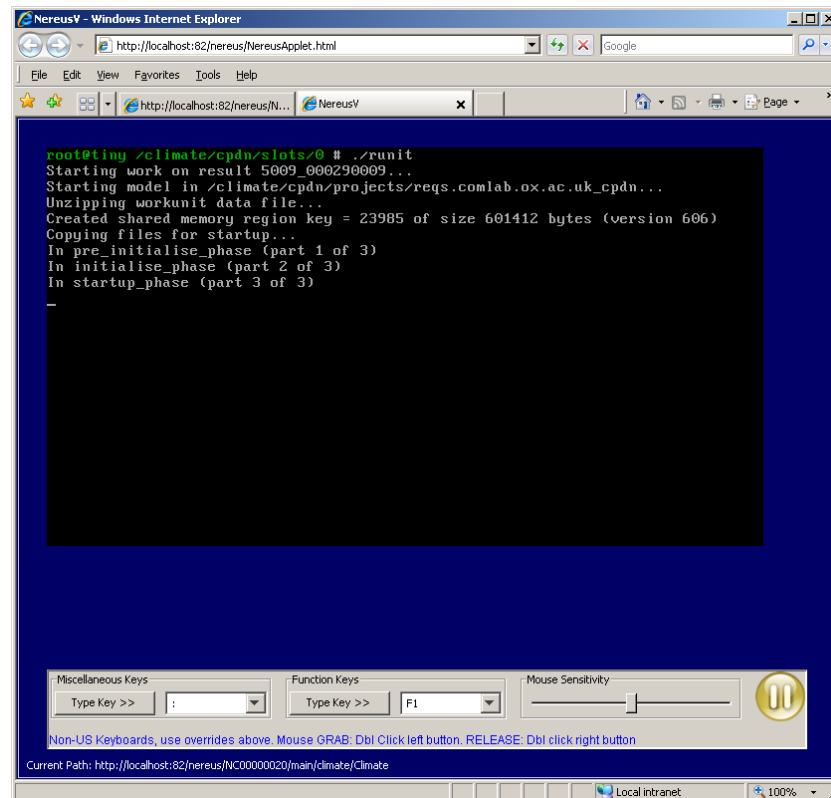
NereusV and JPC

- > Collect idle computing resources together with NereusV
- > Manage these with NereusCloud
- > Deploy “hibernated” virtual machines in vast numbers
- > Thousands of computers doing your processing...
 - On demand
 - Without approval processes / vetting procedures
 - Very cheaply
 - Environmentally friendly



NereusV and JPC Demo

- > Go to our local Nereus Server
- > Watch as we remotely install JPC...
 - Climate Prediction.net code inside TTY Linux













```
root@tiny /climate/cpdn/slots/0 # ./runit
Starting work on result 5009_000290009...
Starting model in /climate/cpdn/projects/reqs.comlab.ox.ac.uk_cpdn...
Unzipping workunit data file...
Created shared memory region key = 23985 of size 601412 bytes (version 606)
Copying files for startup...
In pre_initialise_phase (part 1 of 3)
In initialise_phase (part 2 of 3)
In startup_phase (part 3 of 3)
```

Miscellaneous Keys: Type Key >> :
Function Keys: Type Key >> F1
Mouse Sensitivity: [Slider]
Non-US Keyboards, use overrides above. Mouse GRAB: Dbl Click left button. RELEASE: Dbl click right button
Current Path: http://localhost:82/nereus/NC00000020/main/Climate



So many things to do....



Participation Request

- > Can you help build the NereusV network for research?
- > NereusV is the most cost effective source of CPU power
 - Much research to be done – no \$\$\$ for supercomputers!
 - Benefits to all – results published
- > Do you know a problem NereusV could solve?



Computing of, by, and for the Community

<http://www-nereus.physics.ox.ac.uk>



JavaOneSM

Thank You



Rhys Newman & Ian Preston

rhys.newman@physics.ox.ac.uk

i.preston1@physics.ox.ac.uk

TS-5136