



JavaOne™

java.sun.com/javaone

Everything Java™ Technology... but Better and Faster: The Evolution of JPC

Rhys Newman

Dept Physics, Oxford University and GSA Capital

Chris Dennis

Department of Physics Oxford University

TS-5180



DEMO: JPC In Action...



- Run DOS and Linux virtual machines inside a Java Virtual Machine
 - Try Mario Bros, Pac Man and more classics ... all from within an applet in a browser!
 - Visit <http://www-jpc.physics.ox.ac.uk>

DEMO: Want to see more...?

- ▶ Latest JPC implements the x86 *Virtual8086 mode* which lets it run.....



JPC is an *Emulator* not a *Virtualizer*

> A virtualizer

- Uses hardware support to run *guest* software (including an OS) in isolation to the real *host* operating system
- Will only work when guest software (or OS) is built for the actual hardware being used (x86 on x86, not x86 on Sparc).
- Can achieve near native speeds for the guest software



> An emulator

- Models all hardware in software
- No explicit dependency on underlying hardware – **very flexible**
- Can suffer significant speed penalty



DEMO: Emulation can be *very* flexible...

- ▶ JPC can boot unmodified DOS on a mobile phone
 - Native x86 OS software on an ARM11 system



How Does JPC Do it?

- We read the technical specifications of the x86 PC hardware
- We implement the data flow and structures in Java application environment to mimic these
- We correct these implementations where the specs are wrong!
 - Loads of low-level debugging at “virtual” assembly level
 - We use existing (open source) examples to check out how things work
- Simple so far.....?

The hardware inside an x86 PC:

CD-ROM Drive



Network

Interrupt Controller

PCI Host Bridge

PCI ISA Bridge

PS/2 Interface



Keyboard

DMA Controller

Interval Timer

PCI Bus

VGA Graphics

Processor

VGA BIOS

Mouse



Real Time Clock



MMU

Floppy Drive




BIOS

IDE Interface

Memory

Floppy Controller

Hard Disk Drive



Emulation Speed

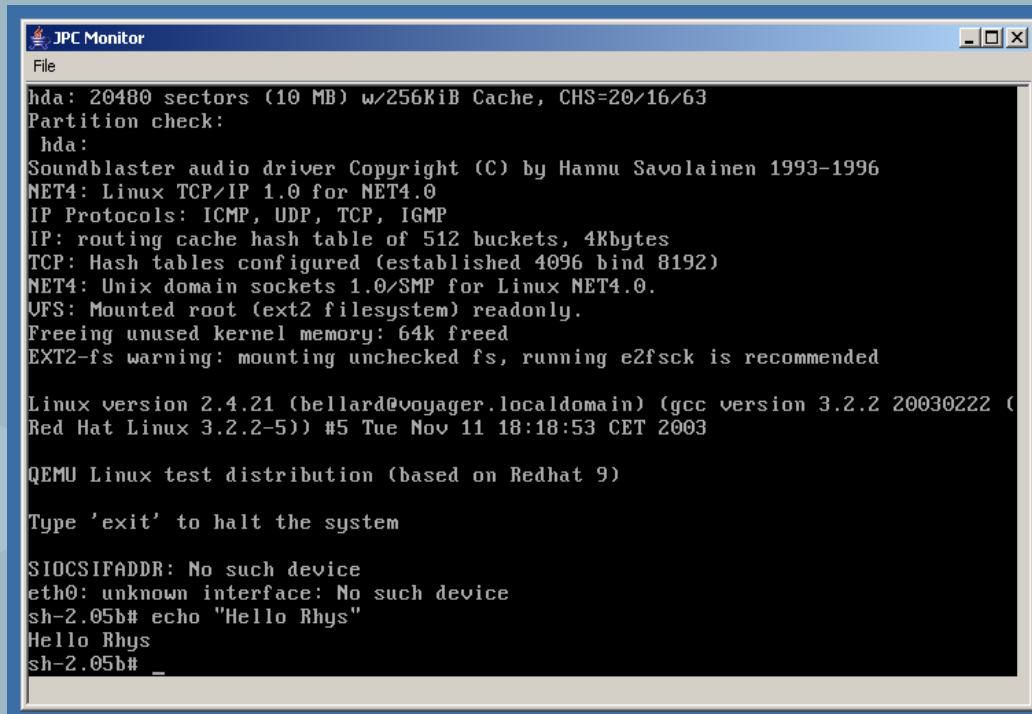
- Doing this the simple naïve way results in a really slow emulator
 - Suffers the double penalty of emulation and JVM overhead (resulting from speed implications of structured OO design)
 - For a much simpler x86 emulator, which is just like this see

IRUC2010 **DIOSCURI**

- But JPC aims to be much faster...
 - Exploit the fact that x86 code is not often changed once loaded
 - Decode once and run many (ratio of 1:100000 is typical)
 - Build classes on the fly and load into the Java Virtual Machine (JVM™ machine) using a runtime classloader – let hotspot go for it....

DEMO: Dynamic Compilation

The only way to run a modern 32bit OS like Linux!



```

JPC Monitor
File
hda: 20480 sectors (10 MB) w/256KiB Cache, CHS=20/16/63
Partition check:
  hda:
Soundblaster audio driver Copyright (C) by Hannu Savolainen 1993-1996
NET4: Linux TCP/IP 1.0 for NET4.0
IP Protocols: ICMP, UDP, TCP, IGMP
IP: routing cache hash table of 512 buckets, 4Kbytes
TCP: Hash tables configured (established 4096 bind 8192)
NET4: Unix domain sockets 1.0/SMP for Linux NET4.0.
VFS: Mounted root (ext2 filesystem) readonly.
Freeing unused kernel memory: 64k freed
EXT2-fs warning: mounting unchecked fs, running e2fsck is recommended

Linux version 2.4.21 (bellard@voyager.localdomain) (gcc version 3.2.2 20030222 (
Red Hat Linux 3.2.2-5)) #5 Tue Nov 11 18:18:53 CET 2003

QEMU Linux test distribution (based on Redhat 9)

Type 'exit' to halt the system

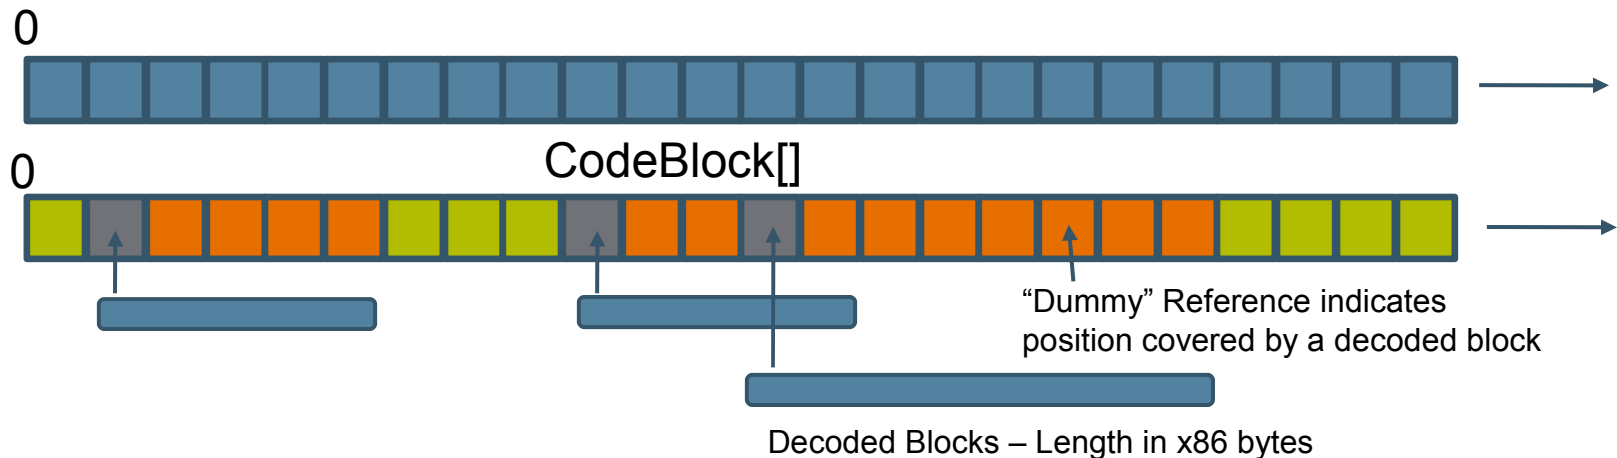
SIOCSIFADDR: No such device
eth0: unknown interface: No such device
sh-2.05b# echo "Hello Rhys"
Hello Rhys
sh-2.05b# _

```

...and we get hardware level checkpointing for free!

How to get the speed

- Instruction pointer jumps to memory address
- Decode x86 instructions up to next jump
- Place object in shadow CodeBlock array which executes the instructions
 - This object can simply interpret x86 in a basic manner
 - Group of x86 can be bunched up, examined and a class file built on the fly to imitate the instructions – *it's our own bytecode compiler*
 - Class file then loaded via a special classloader ready for fast execution next time



The 2007 JavaOne Event Compiler

➤ Let's Be Honest

- The compiler shown at the 2007 JavaOne conference was dumb


➤ It needed rewriting...

- We need smaller classes...
- We need shorter startup code for each block....
- We need more efficient code
 - Dead/redundant code elimination
 - Other standard compiler optimisations
- We need to chain blocks together....

➤ Gentlemen, we can rebuild him. We have the technology

We Form a Directed Acyclic Graph

```
load reg0, eax;  
load reg1, ebx;  
  
add reg0, reg1;  
  
load reg1, ecx;  
sub reg0, reg1;  
  
store eax;
```



Then rinse and repeat for every affected variable.

Then We Generate The Code...

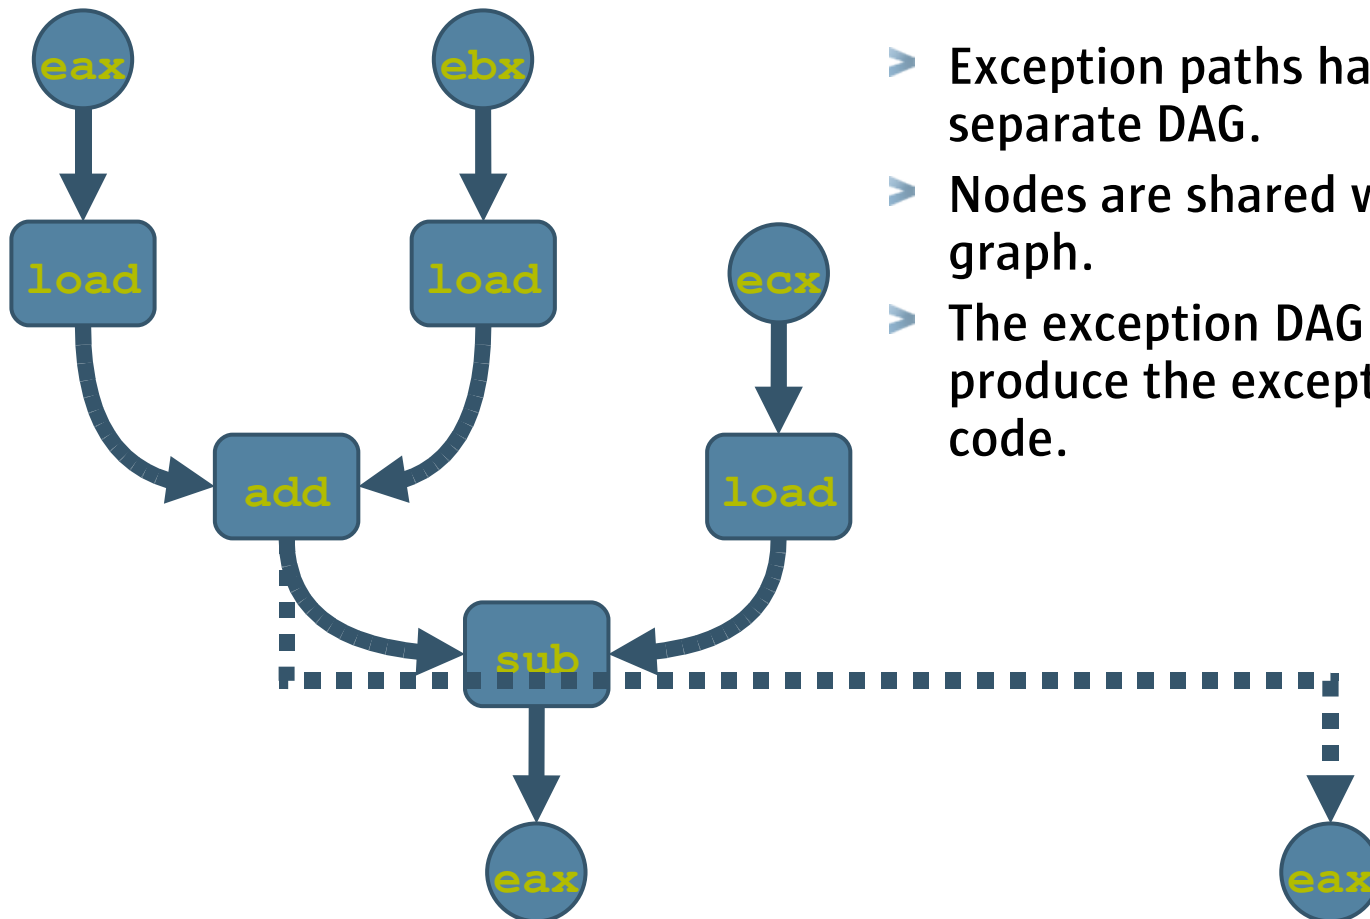
- Code generation from a DAG is simply performed as a depth first traversal of the graph.
- For optimal traversal we should always follow the deepest path at each junction.
 - In a register based machine this optimally reduces the amount of ‘register juggling’
 - In a stack based machine it will give us the minimal stack depth usage.
- Naturally reject unused code – unused code is orphaned from the graph’s final sink
- By tracking multiple node visits we can also eliminate duplicate code.

Damn Those Exceptions...

- In real mode, this scheme works fine as processor (hardware level) exceptions are rare and serious
- In protected mode, the MMU paging system means *every* memory access potentially throws a page fault
 - Not all exceptions are bad, they're just exceptional!
 - In modern machines, this is common
 - After the page fault is handled (memory swapped back in), execution must continue from where it left off
- If this happens within a compiled block of bytecode:
 - Map processor exceptions to Java exceptions and use the JVM exception handling mechanism

Again, again!

Yet more Directed Acyclic Graphs



- > Exception paths have their own separate DAG.
- > Nodes are shared with the main graph.
- > The exception DAG is traversed to produce the exception handler code.

Compiling in Protected Mode

- Use the Java class files exception table to catch page faults
 - The memory object throws a Page Fault exception as required
 - The exception table is populated with special code during compilation to save out processor/memory state up to the point of the fault
 - The fault is then handled as normal
 - The instruction pointer returns to the required location – now part way through the codeblock being executed previously
 - A new codeblock is constructed for that location – most likely JPC will simply interpret the x86 code to the end of the block
 - No problem, next time around the entire block will probably execute in compiled mode as desired!

Better than he was before.

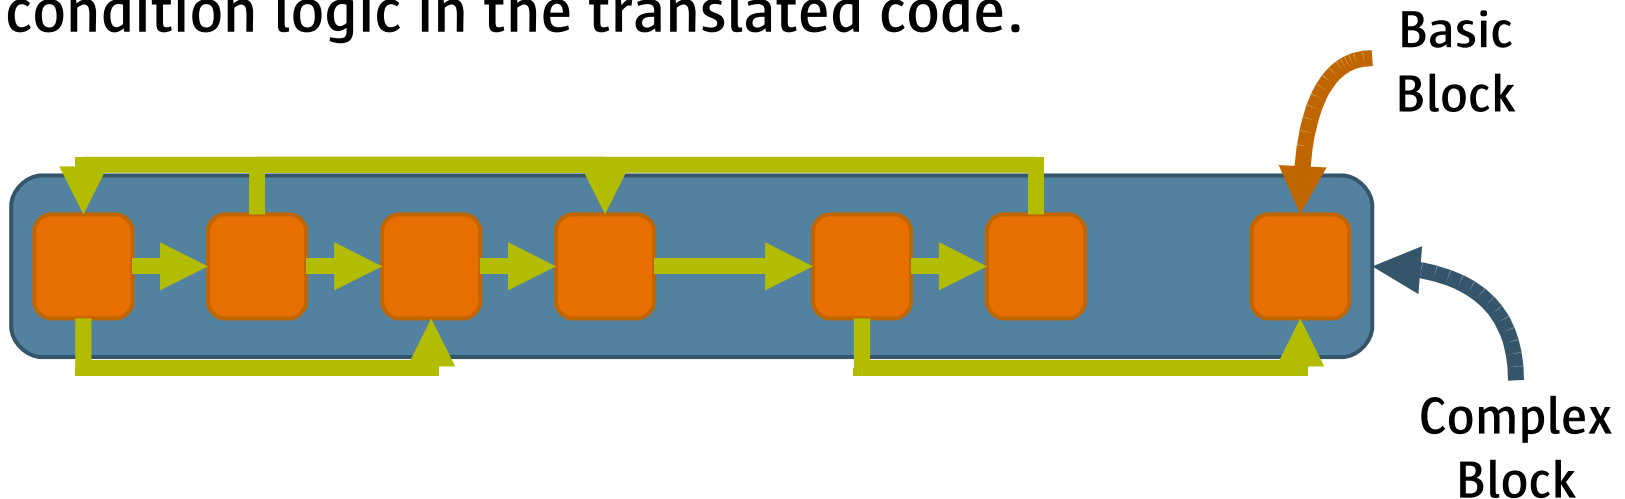
Better, stronger, faster.

- We can now compile blocks in all modes (Real, Protected, Virtual 8086...)
- With full exception handling.
- Redundant code removal.
- Duplicate code removal.
- But this is not the end...

Compilers:


The Next Generation – Invocation Chaining

- Chain invocation of basic blocks by stitching together with appropriate control flow logic.
- Allows the JVM to gain knowledge of the higher level loop and condition logic in the translated code.



- JITed environments like Java Hotspot™ performance engine then have far more scope for optimisation – SPEED!

Summary of JPC Progress



But what is JPC *really* for?



- If you need to isolate guest software from the host, use a virtualiser and get virtually no speed penalty
 - Virtualisation is well understood and can be done for free
 - Major vendors are increasingly supporting virtualisation – both at the hardware level (AMD Pacifica, Intel VT) and at the vendor level

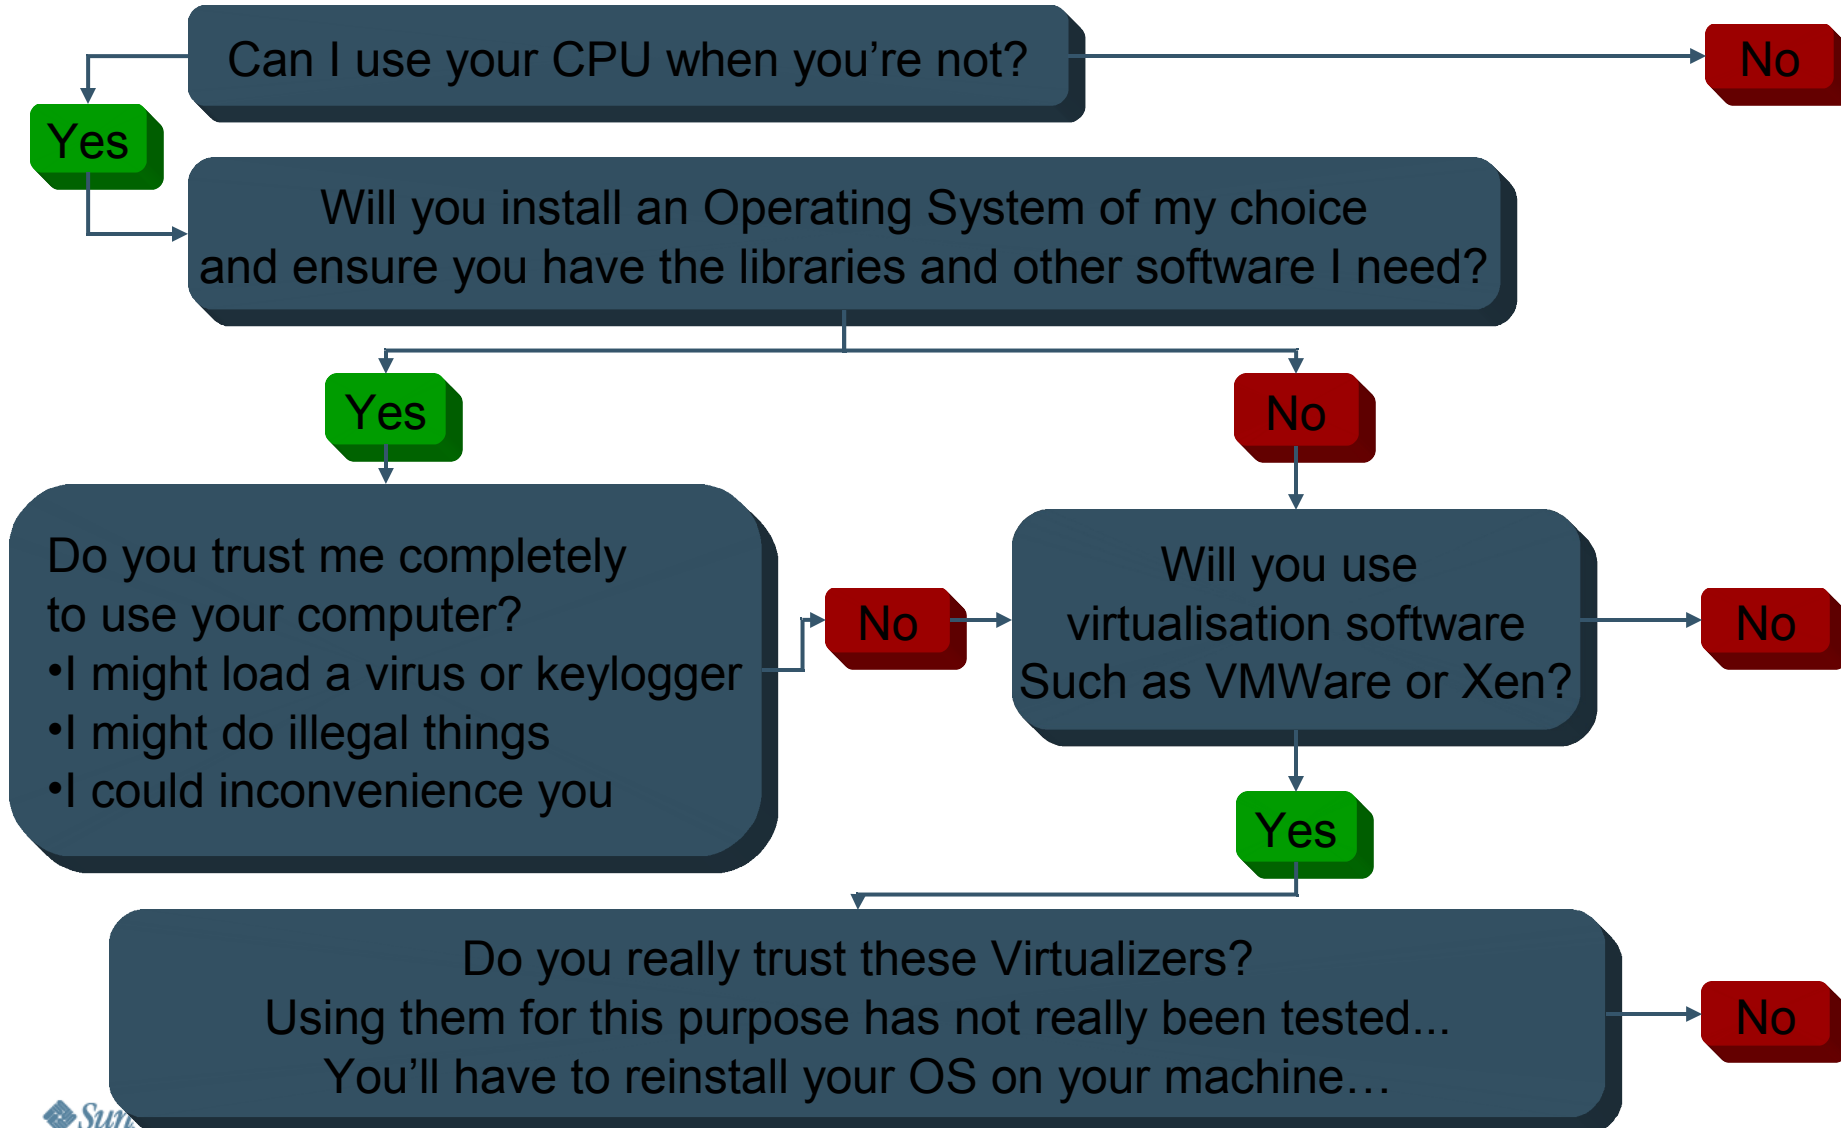
- **But** what if you need so many machines you can't possibly buy them or find a datacentre big enough?
 - There are over 1 billion PCs in the world
 - If you could “borrow” these when idle you could get an unimaginably large amount of CPU power
 - There *are* projects which need this level of processing....

climateprediction.net

SETI HOME
Needs your Help
Donate to SETI@home

Folding@home
distributed computing

A Typical Conversation with a Donor....



Security for Massively Parallel Computations

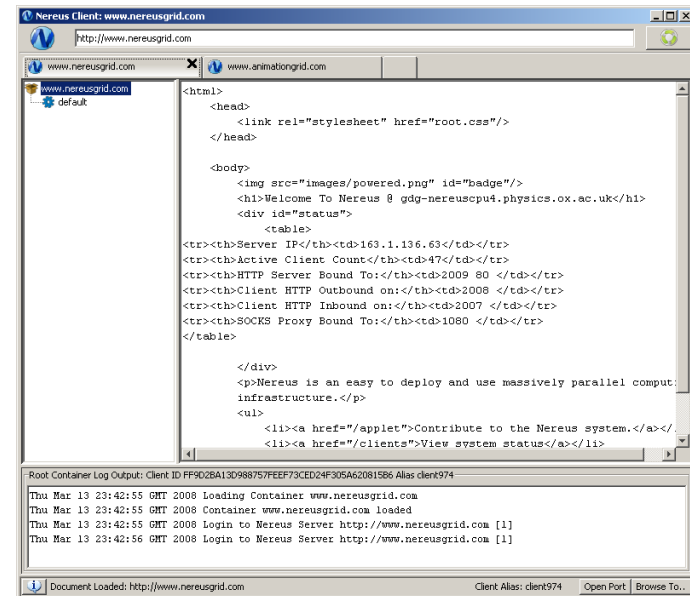
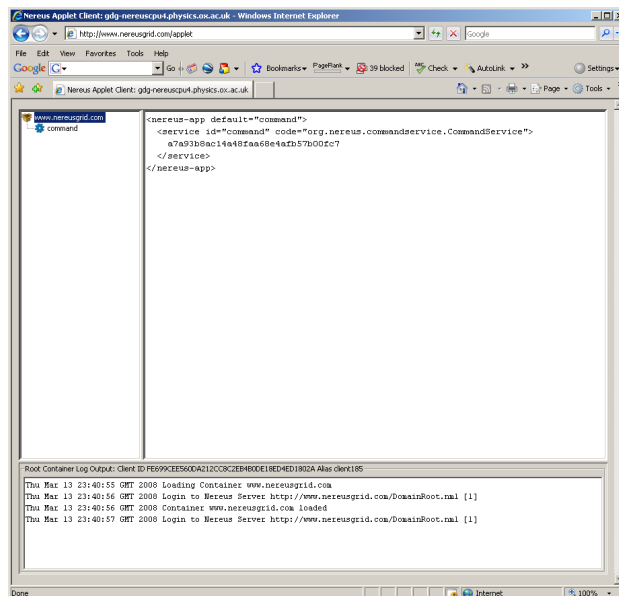
- To harness all these computers you need *absolutely bulletproof* security
 - You can't vet the code before it executes to ensure it isn't stealing private data or doing other naughty things
- You need to stay inside a Sandbox, just like the Applet sandbox
 - has run unvetted code on people's machines for over 10 years
- So you need JPC and a distribution management system for Java technology-based code...

Enter Nereus.....



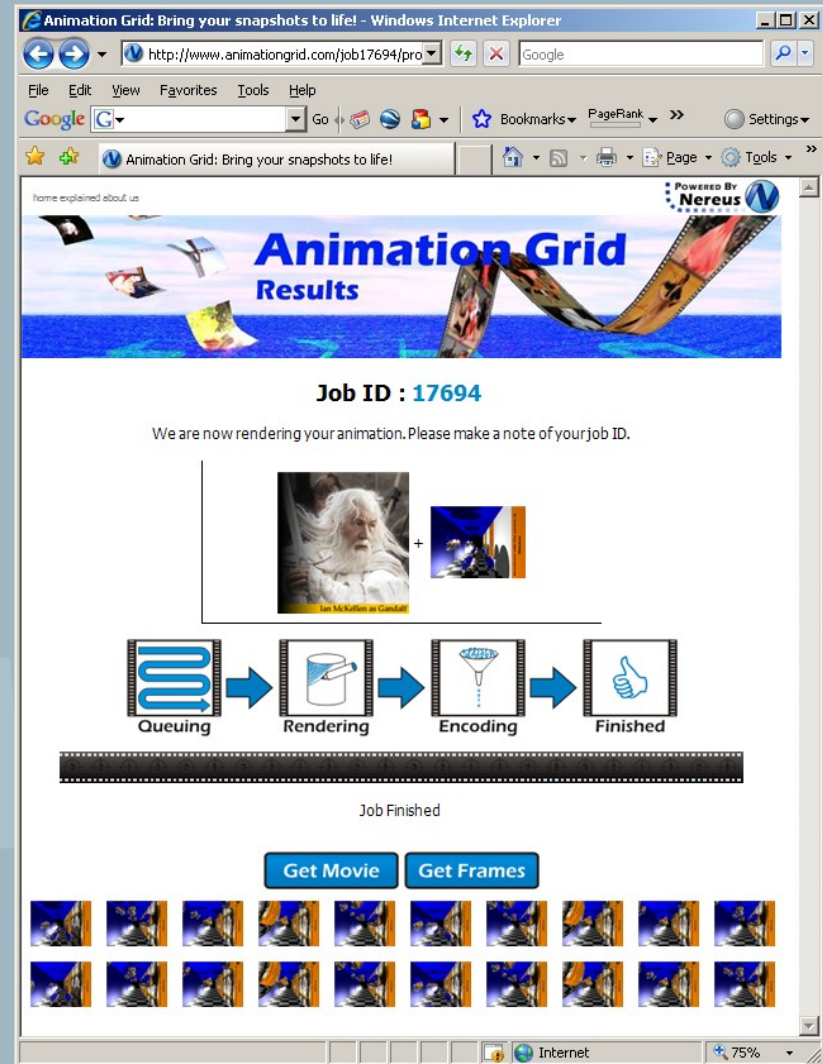
What is Nereus?

- Nereus is an open source pure Java technology infrastructure for MPC
 - Already in use for research, tested on >500 nodes
- To contribute to a Nereus installation either
 - Browse to the installation's web page and click on a link – you'll get the Nereus Applet
 - Download the Nereus Client Application and browse to the installation's root (a normal URL)

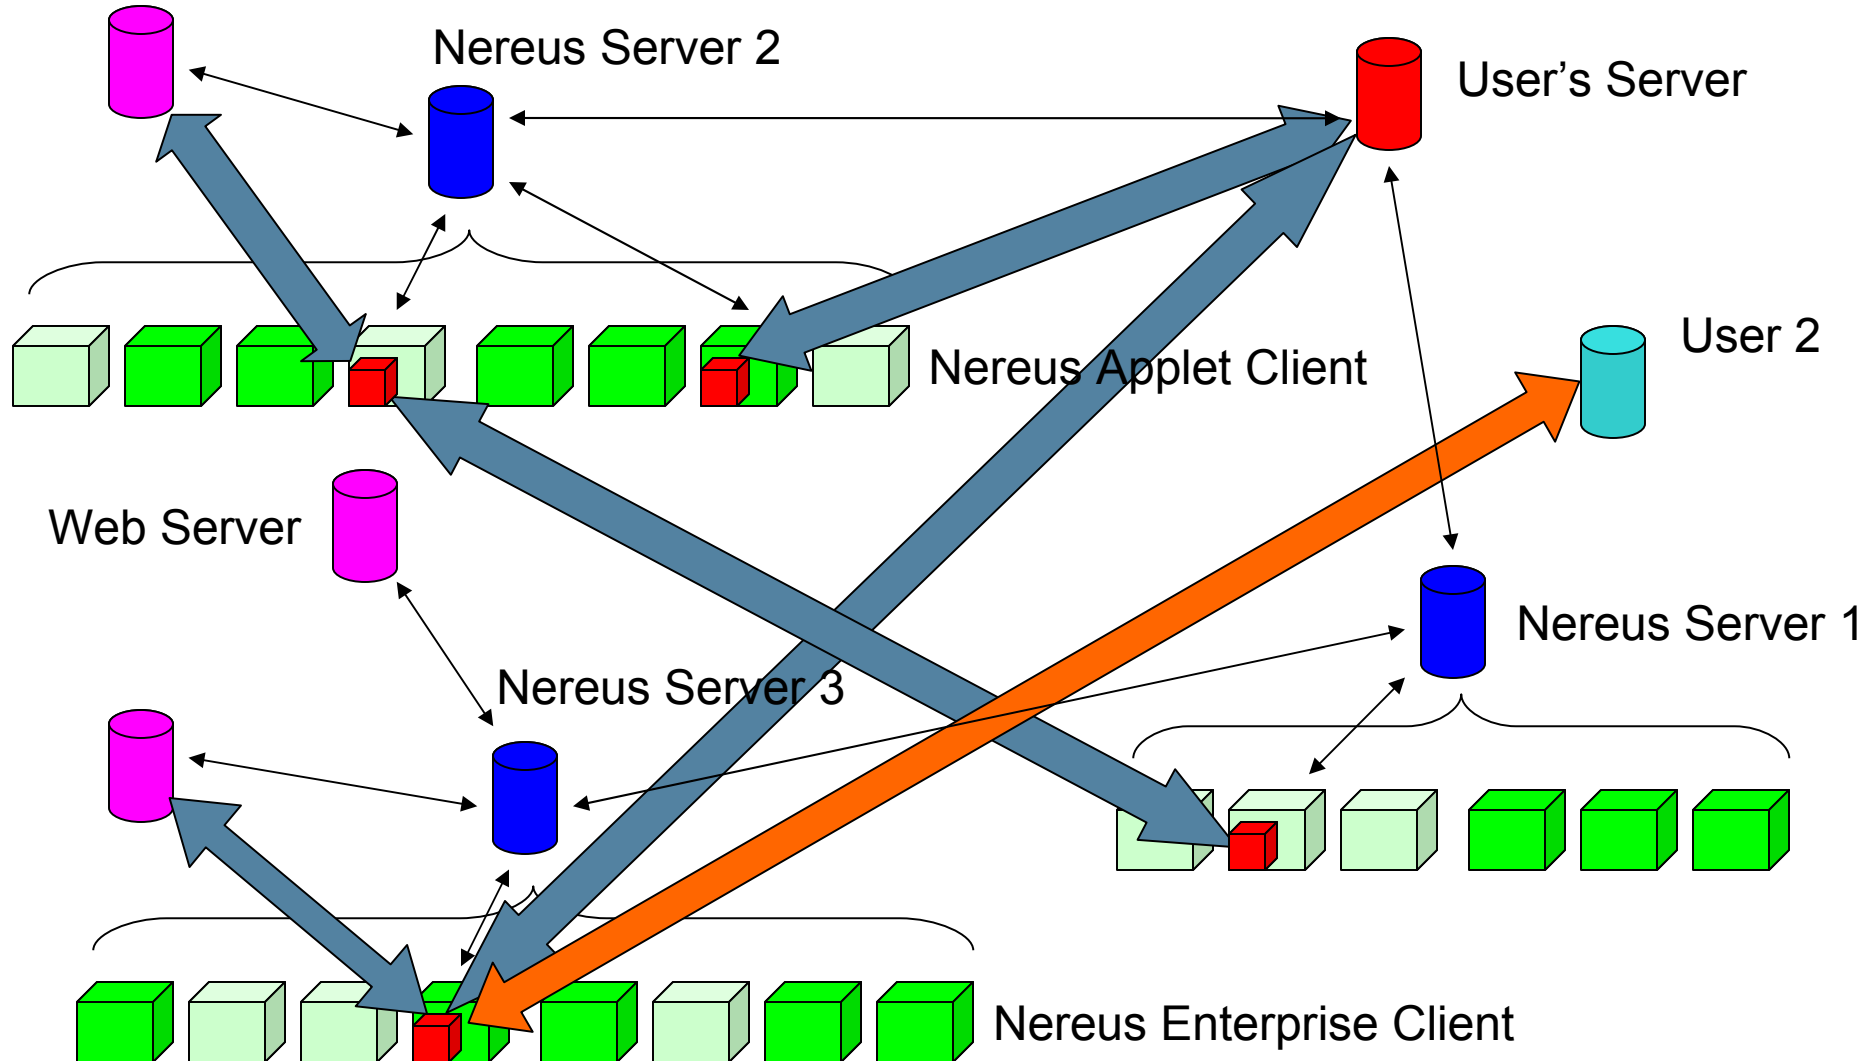


Demo: AnimationGrid

Using “Art of Illusion” Java technology-based renderer to render separate frames on different clients



Nereus System Design:

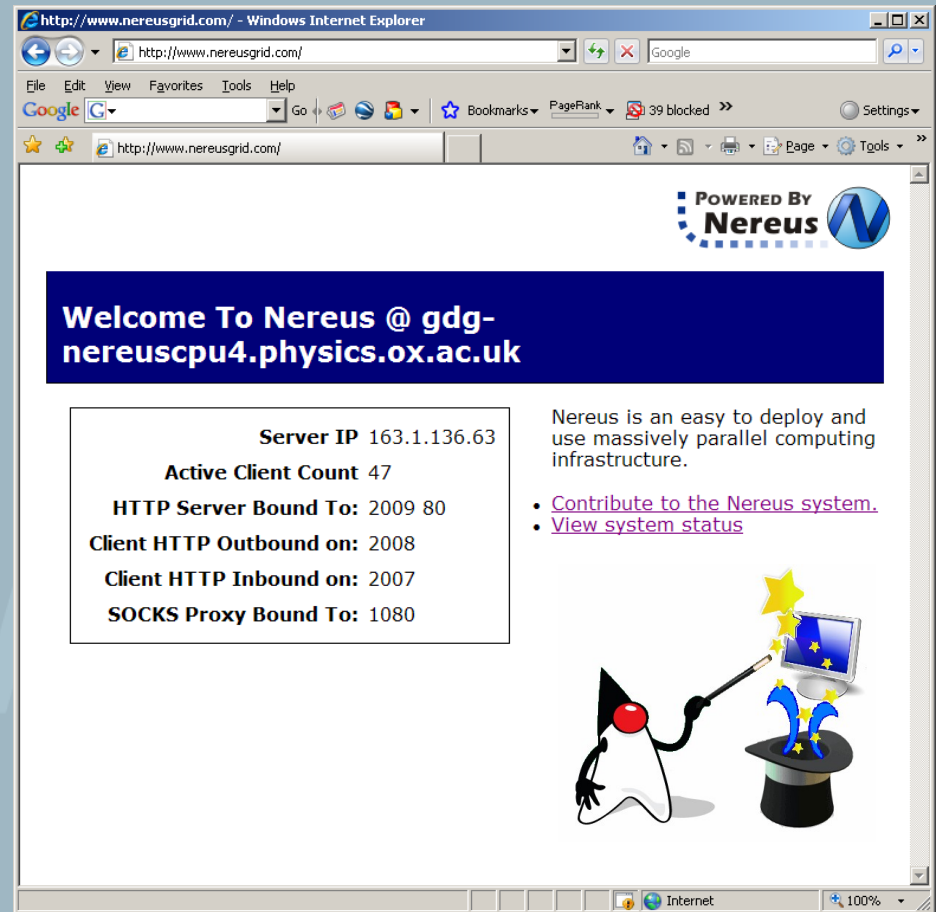


Bringing it all together...

- Use Nereus to manage the distributed resources
- Use JPC to enable arbitrary x86 code (include your own OS) to execute within a *double insulated* sandbox
- Together you get millions of machines able to run your processing
 - On demand
 - Without prior approval processes/vetting procedures
 - Very cheaply – you're renting an existing computer
 - Environmentally friendly – the host machine would probably be on anyway and you're saving energy by not buying more machines in huge datacentres to do your processing!

Demo: JPC on Nereus

How to remotely instantiate JPC instances on Nereus resources – and run native x86 software in an ad-hoc rented machine!



Demo: “Can I have a Volunteer?”

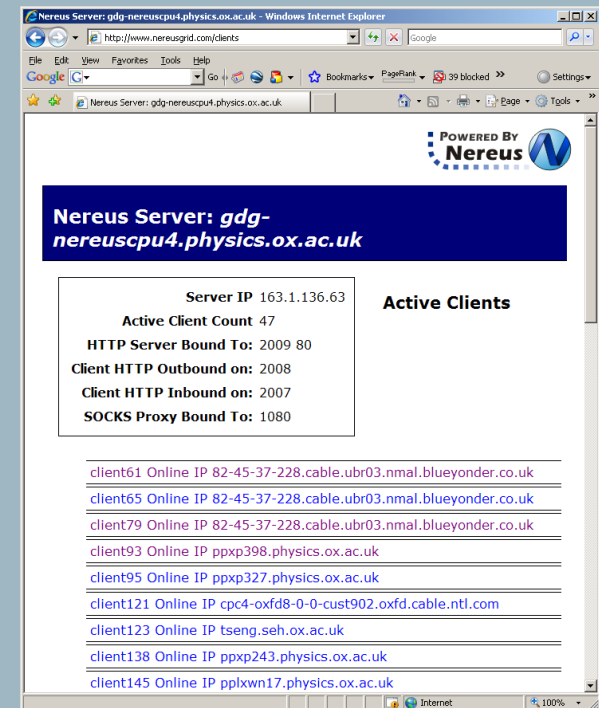


Please open your laptops and browse to
<http://www.nereusgrid.com/>

Note your “clientID”

We will now start JPC on your machine (remotely from the podium), and boot Linux inside your JPC/Nereus applet.

We’re in control, of our “machine”, and you retain control of yours!



Nereus Server: gdg-nereuscpu4.physics.ox.ac.uk

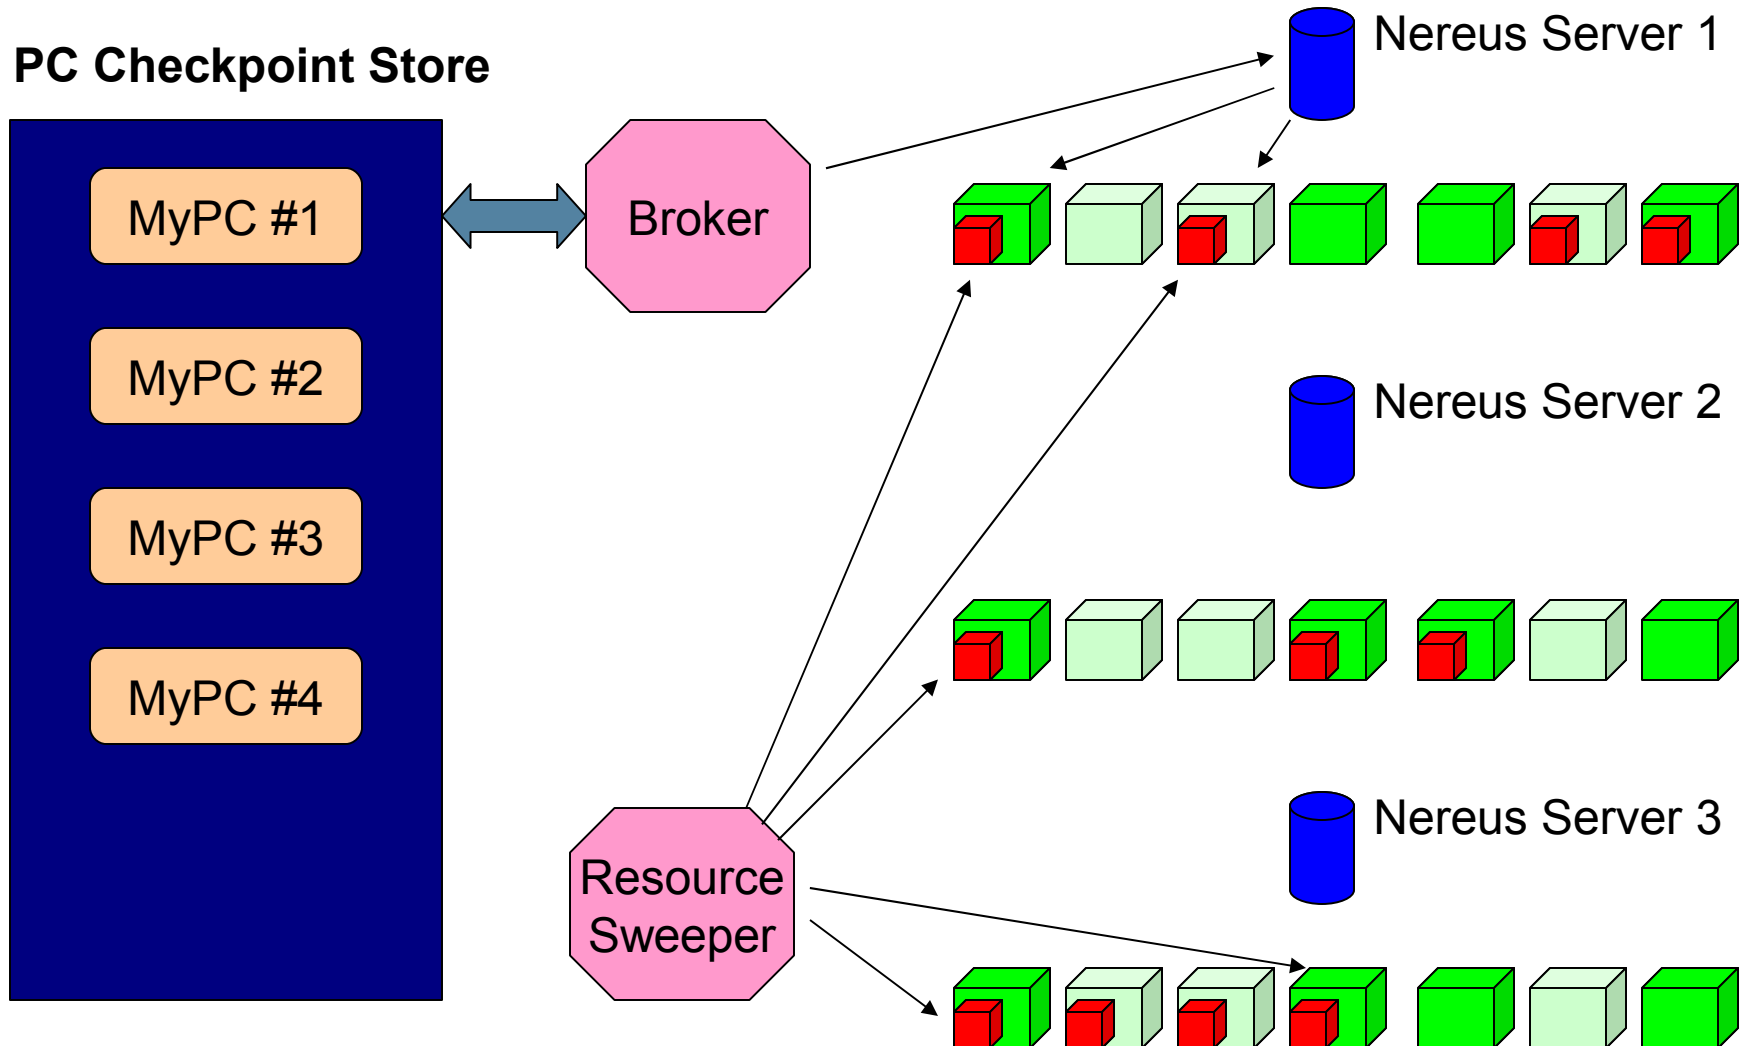
Server IP 163.1.136.63
Active Client Count 47
HTTP Server Bound To: 2009 80
Client HTTP Outbound on: 2008
Client HTTP Inbound on: 2007
SOCKS Proxy Bound To: 1080

Active Clients

- client61 Online IP 82-45-37-228.cable.ubr03.nmal.blueyonder.co.uk
- client65 Online IP 82-45-37-228.cable.ubr03.nmal.blueyonder.co.uk
- client79 Online IP 82-45-37-228.cable.ubr03.nmal.blueyonder.co.uk
- client93 Online IP ppxp398.physics.ox.ac.uk
- client95 Online IP ppxp327.physics.ox.ac.uk
- client121 Online IP cpc4-oxfd8-0-0-cust902.oxfd.cable.ntl.com
- client123 Online IP tseng.seh.ox.ac.uk
- client138 Online IP ppxp243.physics.ox.ac.uk
- client145 Online IP pplxwn17.physics.ox.ac.uk

Global MPC Infrastructure

PC Checkpoint Store



Conclusions

- With over 1 Billion PCs in the world most of the time, there is a vast pool of CPU power going to waste every day
- JPC, combined with Nereus, offers a potential means to release this power by addressing the critical issue of security
- For MPC, where millions machines are needed, there is no practical alternative to this type of technology
- JPC and Nereus offer this opportunity for the Java platform
- With the next level of compilers JPC should be able to reach 50% native speed – well past the tipping point for this global potential
- Download JPC/Nereus and help us out! All our technology is open source GPLv2

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

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

THANK YOU



Rhys Newman

Dept Physics, Oxford University and GSA Capital

Chris Dennis

Department of Physics Oxford University

TS-5180

