cse

Extended OS

THE UNIVERSITY OF NEW SOUTH WALES

cse

Virtual Machines

References:

Smith, J.E.; Ravi Nair; , "The architecture of virtual machines," *Computer* , vol.38, no.5, pp. 32- 38, May 2005

THE UNIVERSITY OF NEW SOUTH WALES

cse

Abstraction & Virtualisation

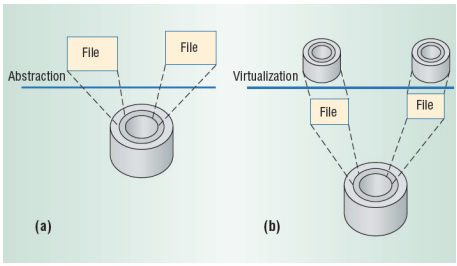

THE UNIVERSITY OF NEW SOUTH WALES

cse

Interface Levels

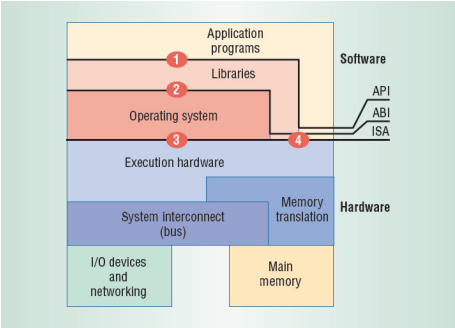

THE UNIVERSITY OF NEW SOUTH WALES

cse

Instruction Set Architecture

- Interface between software and hardware
- Divided between privileged and un-privileged parts

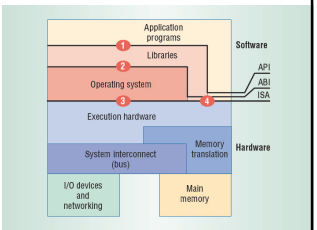

THE UNIVERSITY OF NEW SOUTH WALES

cse

Application Binary Interface

- Interface between programs hardware + OS
- Consists of system call interface + un-privileged ISA

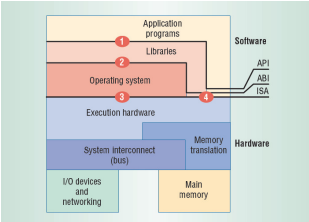

THE UNIVERSITY OF NEW SOUTH WALES

cse

Application Programming Interface

- Interface between programs hardware + OS
- Consists of library calls + un-privileged ISA
  - Syscalls usually called through library.

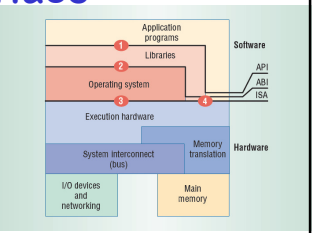

THE UNIVERSITY OF NEW SOUTH WALES

cse

Process versus *System* Virtual Machine

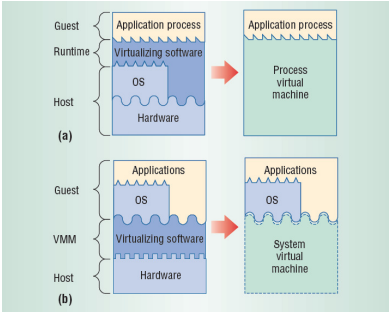

THE UNIVERSITY OF NEW SOUTH WALES

cse

OS is an extended virtual machine

- Multiplexes the "machine" between applications
  - Time sharing, multitasking, batching
- Provided a higher-level machine for
  - Ease of use
  - Portability
  - Efficiency
  - Security
  - Etc….

THE UNIVERSITY OF NEW SOUTH WALES

cse

JAVA – Higher-level Virtual Machine

- write a program once, and run it anywhere
  - Architecture independent
  - Operating System independent
- Language itself was clean, robust, garbage collection
- Program compiled into bytecode
  - Interpreted or just-in-time compiled.
  - Lower than native performance

THE UNIVERSITY OF NEW SOUTH WALES

cse

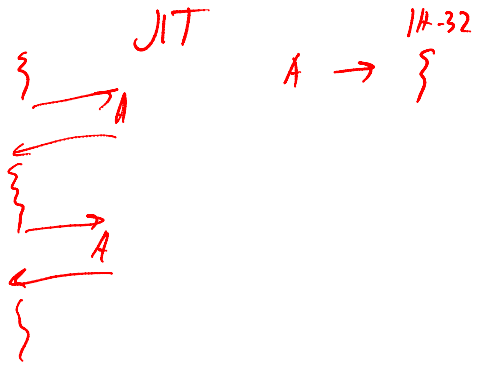

THE UNIVERSITY OF NEW SOUTH WALES

cse

Conventional versus Emulation/Translation

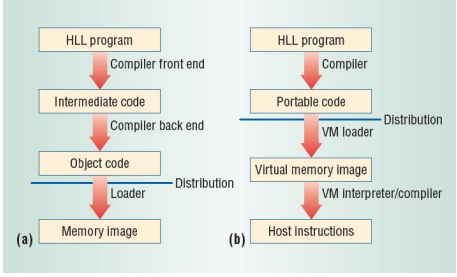

THE UNIVERSITY OF NEW SOUTH WALES

cse

Issues

- Legacy applications
- No isolation nor resource management between applets
- Security
  - Trust JVM implementation? Trust underlying OS?
- Performance compared to native

THE UNIVERSITY OF NEW SOUTH WALES

cse

Is the OS the "right" level of extended machine?

- Security
  - Trust the underlying OS?
- Legacy application and OSs
- Resource management of existing systems suitable for all applications?
- What about activities requiring "root" privileges

THE UNIVERSITY OF NEW SOUTH WALES

cse

Virtual Machine Monitors

- Provide scheduling and resource management
- Extended "machine" is the actual machine interface.

THE UNIVERSITY OF NEW SOUTH WALES

cse

IBM VM/370

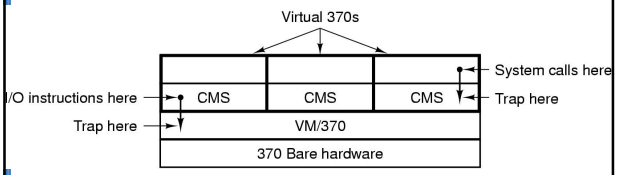

THE UNIVERSITY OF NEW SOUTH WALES

cse

Advantages

- Legacy OSes (and applications)
- Server consolidation
- Concurrent OSes
  - Linux – Windows
  - Primary – Backup
    - High availability
- Test and Development
- Security
  - VMM (hopefully) small and correct
- Performance near bare hardware
  - For some applications

THE UNIVERSITY OF NEW SOUTH WALES

cse

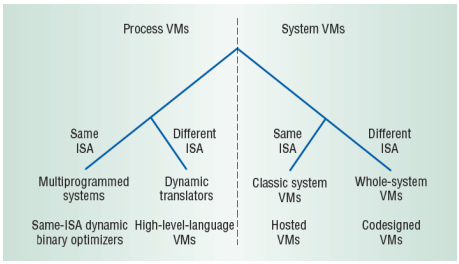

THE UNIVERSITY OF NEW SOUTH WALES

cse

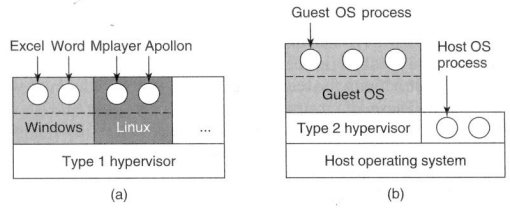

Figure 1-29. (a) A type 1 hypervisor. (b) A type 2 hypervisor.

THE UNIVERSITY OF NEW SOUTH WALES

cse

Virtual R3000???

- Interpret
  - System/161
    - slow
  - JIT dynamic compilation
- Run on the real hardware??

THE UNIVERSITY OF NEW SOUTH WALES

cse

R3000 Virtual Memory Addressing

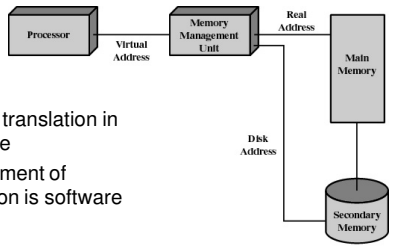

- MMU
  - address translation in hardware
  - management of translation is software

Figure 2.10 Virtual Memory Addressing

THE UNIVERSITY OF NEW SOUTH WALES

21

cse

R3000 Address Space Layout

- kuseg:
  - 2 gigabytes
  - MMU translated
  - Cacheable
  - user-mode and kernel mode accessible

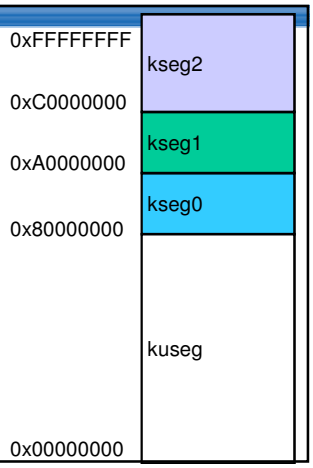

THE UNIVERSITY OF NEW SOUTH WALES

cse

R3000 Address Space Layout

- kseg0:
  - 512 megabytes
  - Fixed translation window to physical memory
    - 0x80000000 - 0x9fffffff virtual = 0x00000000 - 0x1fffffff physical
    - MMU not used
  - Cacheable
  - Only kernel-mode accessible
  - Usually where the kernel code is placed

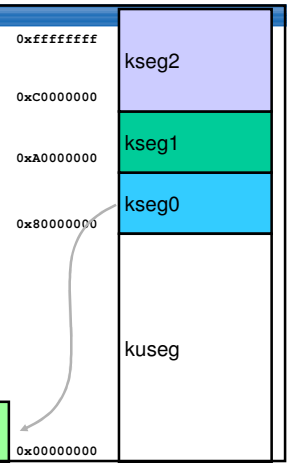

THE UNIVERSITY OF NEW SOUTH WALES

cse

R3000 Address Space Layout

- kseg1:
  - 512 megabytes
  - Fixed translation window to physical memory
    - 0xa0000000 - 0xbfffffff virtual = 0x00000000 - 0x1fffffff physical
    - MMU not used
  - **NOT** cacheable
  - Only kernel-mode accessible
  - Where devices are accessed (and boot ROM)

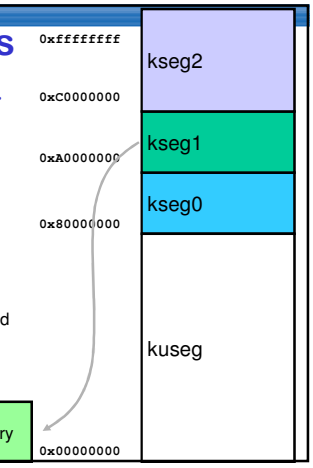

THE UNIVERSITY OF NEW SOUTH WALES

cse

R3000 Address Space Layout

- kseg2:
  - 1024 megabytes
  - MMU translated
  - Cacheable
  - Only kernel-mode accessible

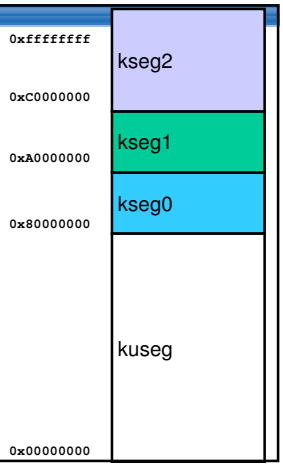

THE UNIVERSITY OF NEW SOUTH WALES

cse

Issues

- Privileged registers (CP0)
- Privileged instructions
- Address Spaces
- Exceptions (including syscalls, interrupts)
- Devices

THE UNIVERSITY OF NEW SOUTH WALES

cse

THE UNIVERSITY OF NEW SOUTH WALES

cse

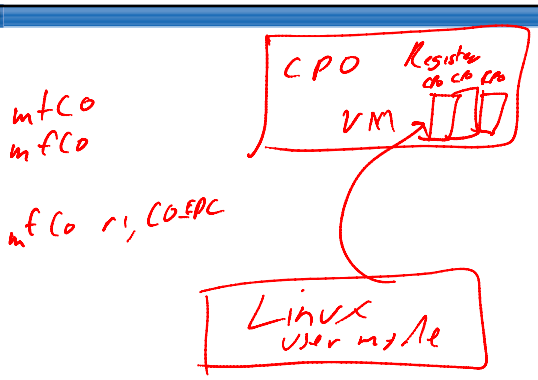

THE UNIVERSITY OF NEW SOUTH WALES

cse

THE UNIVERSITY OF NEW SOUTH WALES

cse

THE UNIVERSITY OF NEW SOUTH WALES

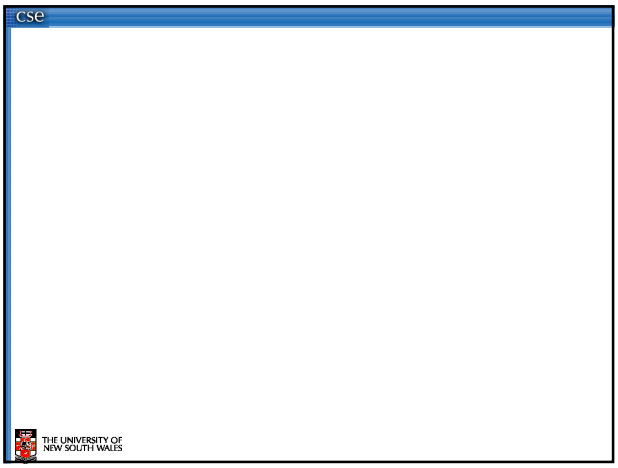
cse
THE UNIVERSITY OF
NEW SOUTH WALES