

Operating Systems Design (CS 423)

Elsa L Gunter
2112 SC, UIUC

<http://www.cs.illinois.edu/class/cs423/>

Based on slides by Roy Campbell, Sam King, and
Andrew S Tanenbaum

4/1/11

1

Simulator

- CPU state --- data structure within your sim.
 - Includes registers, IDT, etc.
- Memory --- malloc
 - Guest physical address 0 == beginning of malloc
- Disk --- file
 - Disk block 1 = file offset 1*(size of disk block)
- Display --- window
- Within simulator, does the software "know" it is being simulated?

4/1/11

2

VMM environment

- Duplicate
 - Virtual machine == physical machine
- Efficient
 - Runs almost as fast as real machine
- Isolated
 - VMM has control over resources
 - What are the resources the VMM controls?
- Which properties does Simulator have?

4/1/11

3

Key observation

- If the virtual ISA == physical ISA
 - CPU can perform simulation loop
- Can we set host PC to address we want to start at and let it run?
- What property are we violating?
- What else goes wrong?

4/1/11

4

Main issues

- Privileged instructions
- Instructions operate on physical state, not virtual state

4/1/11

5

Privileged instructions

- CPU divided into supervisor and user modes
- Part of the CPU ISA only accessible by "supervisor" code
- Allowing guest OS execute these would violate isolation
 - E.g., I/O instructions to write disk blocks
- Solution: run guest OS in user mode
 - CPU user mode: only non-privileged instr
 - CPU supervisor mode: all instr

4/1/11

6


Example: interrupt descriptor table (IDT)

- x86 processor have an IDT register
 - CPU uses to find interrupt service routines
- Set the IDT register using the `lidt` instruction
- VMM must handle all interrupts to maintain control
- Goal: allow the guest OS to set the virtual IDT register without affecting the physical CPU

4/1/11

7


Guest OS priv. inst.



4/1/11

8


Guest OS priv. inst.



4/1/11

9


Guest OS priv. inst.



4/1/11

10

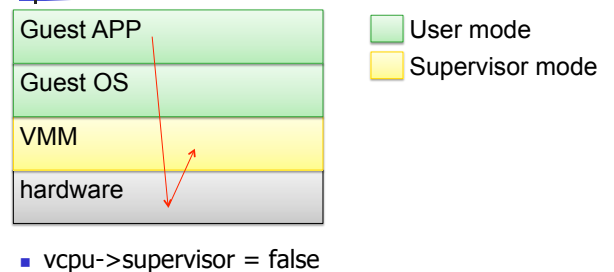
Guest OS priv. inst.



4/1/11

11


Guest OS priv. inst.



4/1/11

12

Guest OS priv. inst.



- vcpu->supervisor = true

4/1/11

13

Sensitive non-privileged instructions

- Privileged instruction
 - Trap when called from CPU user mode
- Sensitive instruction
 - Leaks information about physical state of processor
 - Eg sidt
- Fully virtualizable processor
 - All sensitive instructions are privileged
- What about x86?

4/1/11

14


Sensitive non-privileged instructions

- x86 has a number of instructions that are sensitive, non-privileged instructions
 - Causes physical state of CPU to leak
 - Why is this a problem?
- What property does this violate?

4/1/11

15

Guest OS sensitive non-priv. inst.



- sidt = 0xabcd
- vcpu->idtr = 0x1234
- idtr = 0xabcd

4/1/11

16

Key insight

- Solution: simulate the OS, let user-mode code run natively
 - Simulation is flexible
 - Can interpose on all instructions
 - Problem: simulation is too slow
- Will work if guest OS calls SIDT, what about guest user mode?

4/1/11

17

Simulation flexibility

- Normal simulations flexible, slow
- Can we simulate fast, still flexible?


```
while(1){
    inst = mem[PC]; // fetch
    if(inst == add) { // decode// execute
        reg[inst.reg1] = reg[inst.reg2] + reg[inst.reg3];
        PC++;
    }
    else if ...
} // repeat
```

4/1/11

18

Dynamic binary translation

- Simulate basic block using host instructions.
Basic block?
- Cache translations



4/1/11

19

Performance

```
while(1){  
    inst = mem[PC]; // fetch  
    //decode  
    if(inst == add){  
        //execute  
        ...  
    }  
} // repeat
```

```
while(1){  
    if(!translated(pc)){  
        translate(pc);  
    }  
    jump to pc2tc(pc);  
} // repeat
```

- Why is this faster?
- Are there disadvantages?

4/1/11

20