

**CS162**  
**Operating Systems and**  
**Systems Programming**  
**Lecture 2**

**Introduction to the Process**

**August 31<sup>st</sup>, 2015**

**Prof. John Kubiawicz**


**<http://cs162.eecs.Berkeley.edu>**

*Acknowledgments: Lecture slides are from the Operating Systems course taught by John Kubiawicz at Berkeley, with few minor updates/changes. When slides are obtained from other sources, a reference will be noted on the bottom of that slide, in which case a full list of references is provided on the last slide.*

## Recall: What is an operating system?

---

- Special layer of software that provides application software access to hardware resources
  - Convenient abstraction of complex hardware devices
  - Protected access to shared resources
  - Security and authentication
  - Communication amongst logical entities



# Review: What is an Operating System?

---



- **Referee**

- Manage sharing of resources, Protection, Isolation
  - » Resource allocation, isolation, communication

- **Illusionist**

- Provide clean, easy to use abstractions of physical resources
  - » Infinite memory, dedicated machine
  - » Higher level objects: files, users, messages
  - » Masking limitations, virtualization



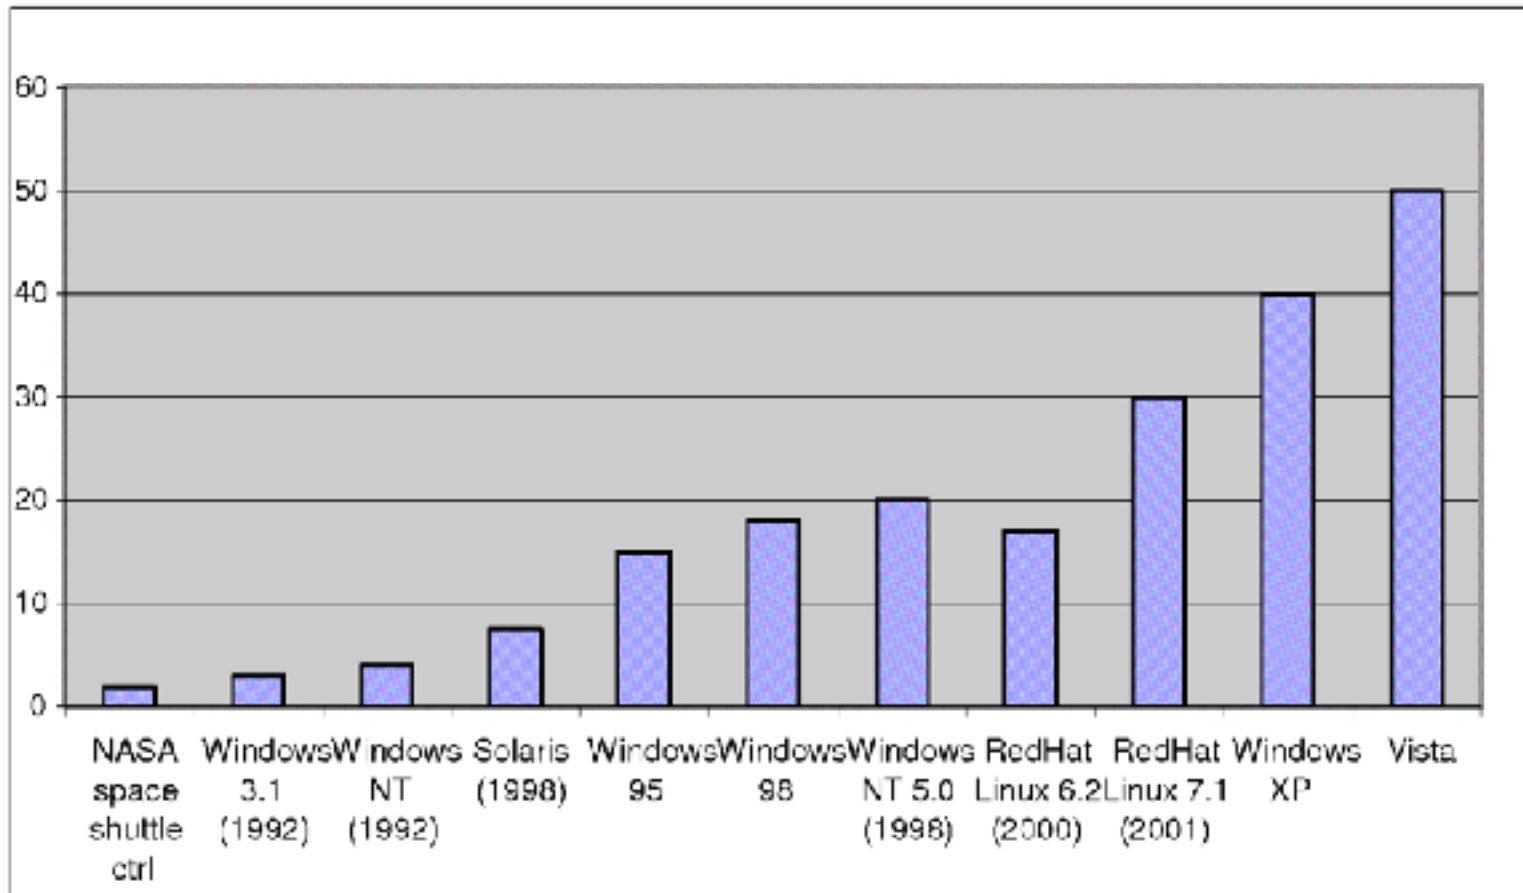
- **Glue**

- Common services
  - » Storage, Window system, Networking
  - » Sharing, Authorization
  - » Look and feel



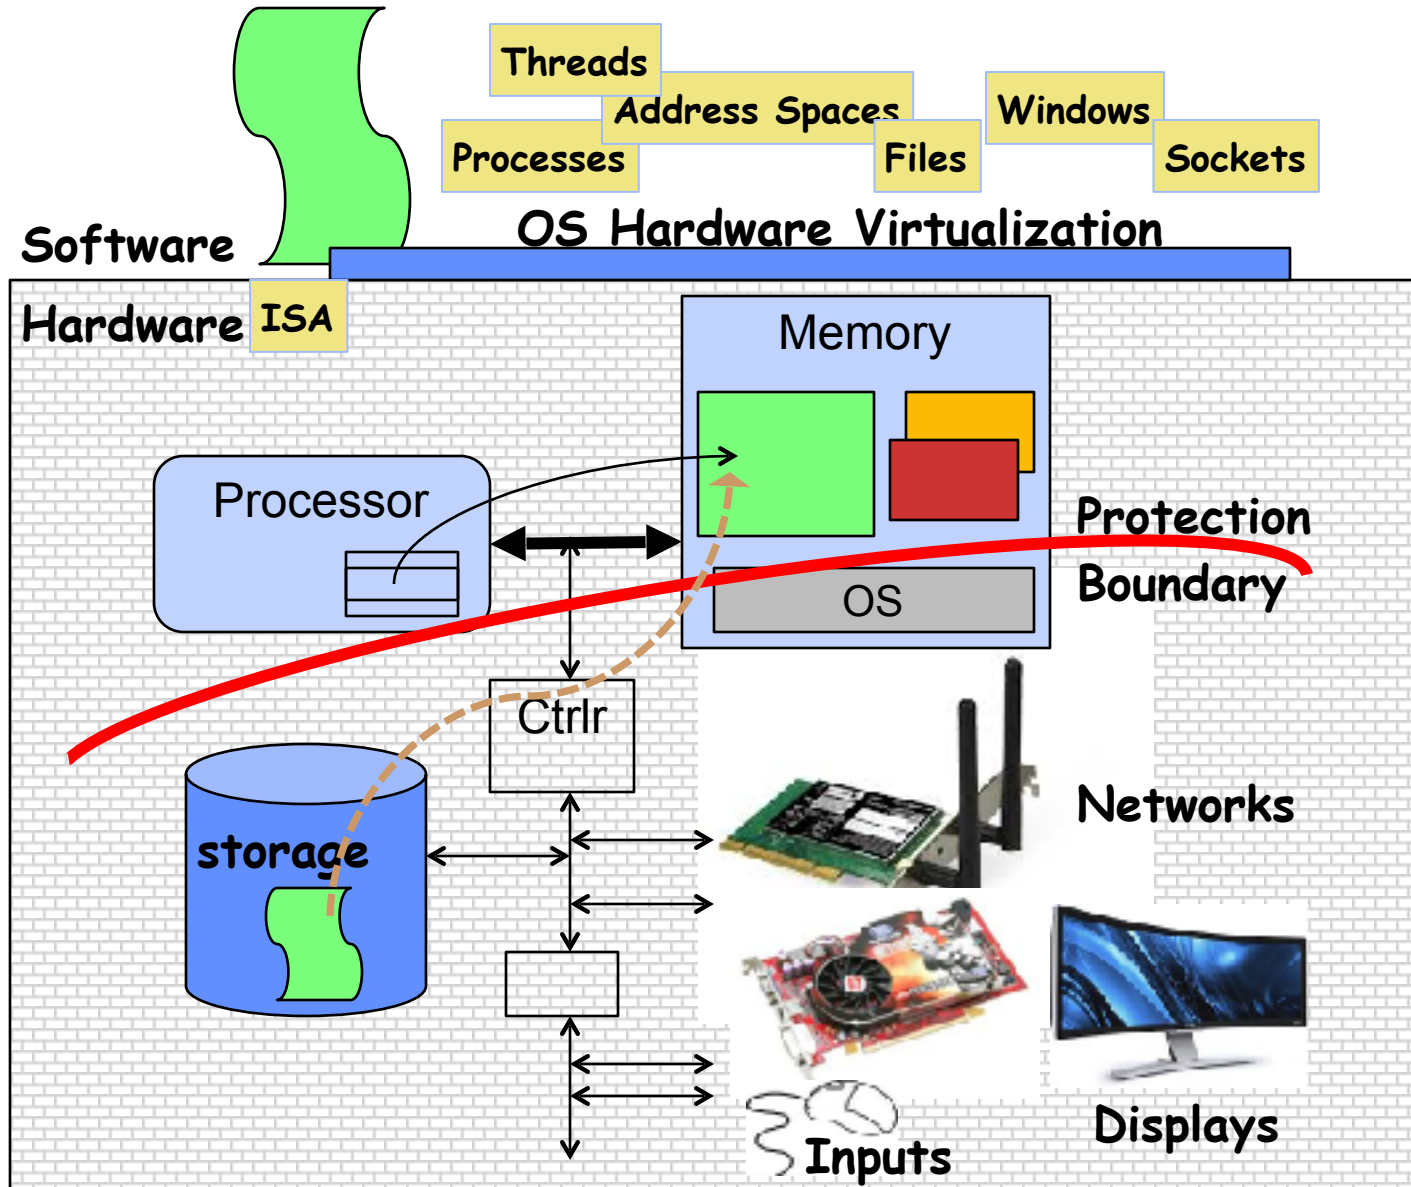
# Review: Increasing Software Complexity

Millions of lines of source code



From MIT's 6.033 course

# Recall: Loading



# Very Brief History of OS

---

- **Several Distinct Phases:**
  - **Hardware Expensive, Humans Cheap**
    - » Eniac, ... Multics
  - **Hardware Cheaper, Humans Expensive**
    - » PCs, Workstations, Rise of GUIs
  - **Hardware Really Cheap, Humans Really Expensive**
    - » Ubiquitous devices, Widespread networking



"I think there is a world market for maybe five computers." -- Thomas Watson, chairman of IBM, 1943

# Very Brief History of OS

---

- **Several Distinct Phases:**
  - **Hardware Expensive, Humans Cheap**
    - » Eniac, ... Multics
  - **Hardware Cheaper, Humans Expensive**
    - » PCs, Workstations, Rise of GUIs
  - **Hardware Really Cheap, Humans Really Expensive**
    - » Ubiquitous devices, Widespread networking




Thomas Watson was often called "the worlds greatest salesman" by the time of his death in 1956

# Very Brief History of OS

---

- **Several Distinct Phases:**
  - **Hardware Expensive, Humans Cheap**
    - » Eniac, ... Multics
  - **Hardware Cheaper, Humans Expensive**
    - » PCs, Workstations, Rise of GUIs
  - **Hardware Really Cheap, Humans Really Expensive**
    - » Ubiquitous devices. Widespread networking





# Very Brief History of OS

---

- **Several Distinct Phases:**
  - **Hardware Expensive, Humans Cheap**
    - » Eniac, ... Multics
  - **Hardware Cheaper, Humans Expensive**
    - » PCs, Workstations, Rise of GUIs
  - **Hardware Really Cheap, Humans Really Expensive**
    - » Ubiquitous devices, Widespread networking
- **Rapid Change in Hardware Leads to changing OS**
  - **Batch ⇒ Multiprogramming ⇒ Timesharing ⇒ Graphical UI ⇒ Ubiquitous Devices**
  - **Gradual Migration of Features into Smaller Machines**
- **Situation today**
  - **Small OS: 100K lines/Large: 10M lines (5M browser!)**
  - **100-1000 people-years**

# OS Archaeology

---

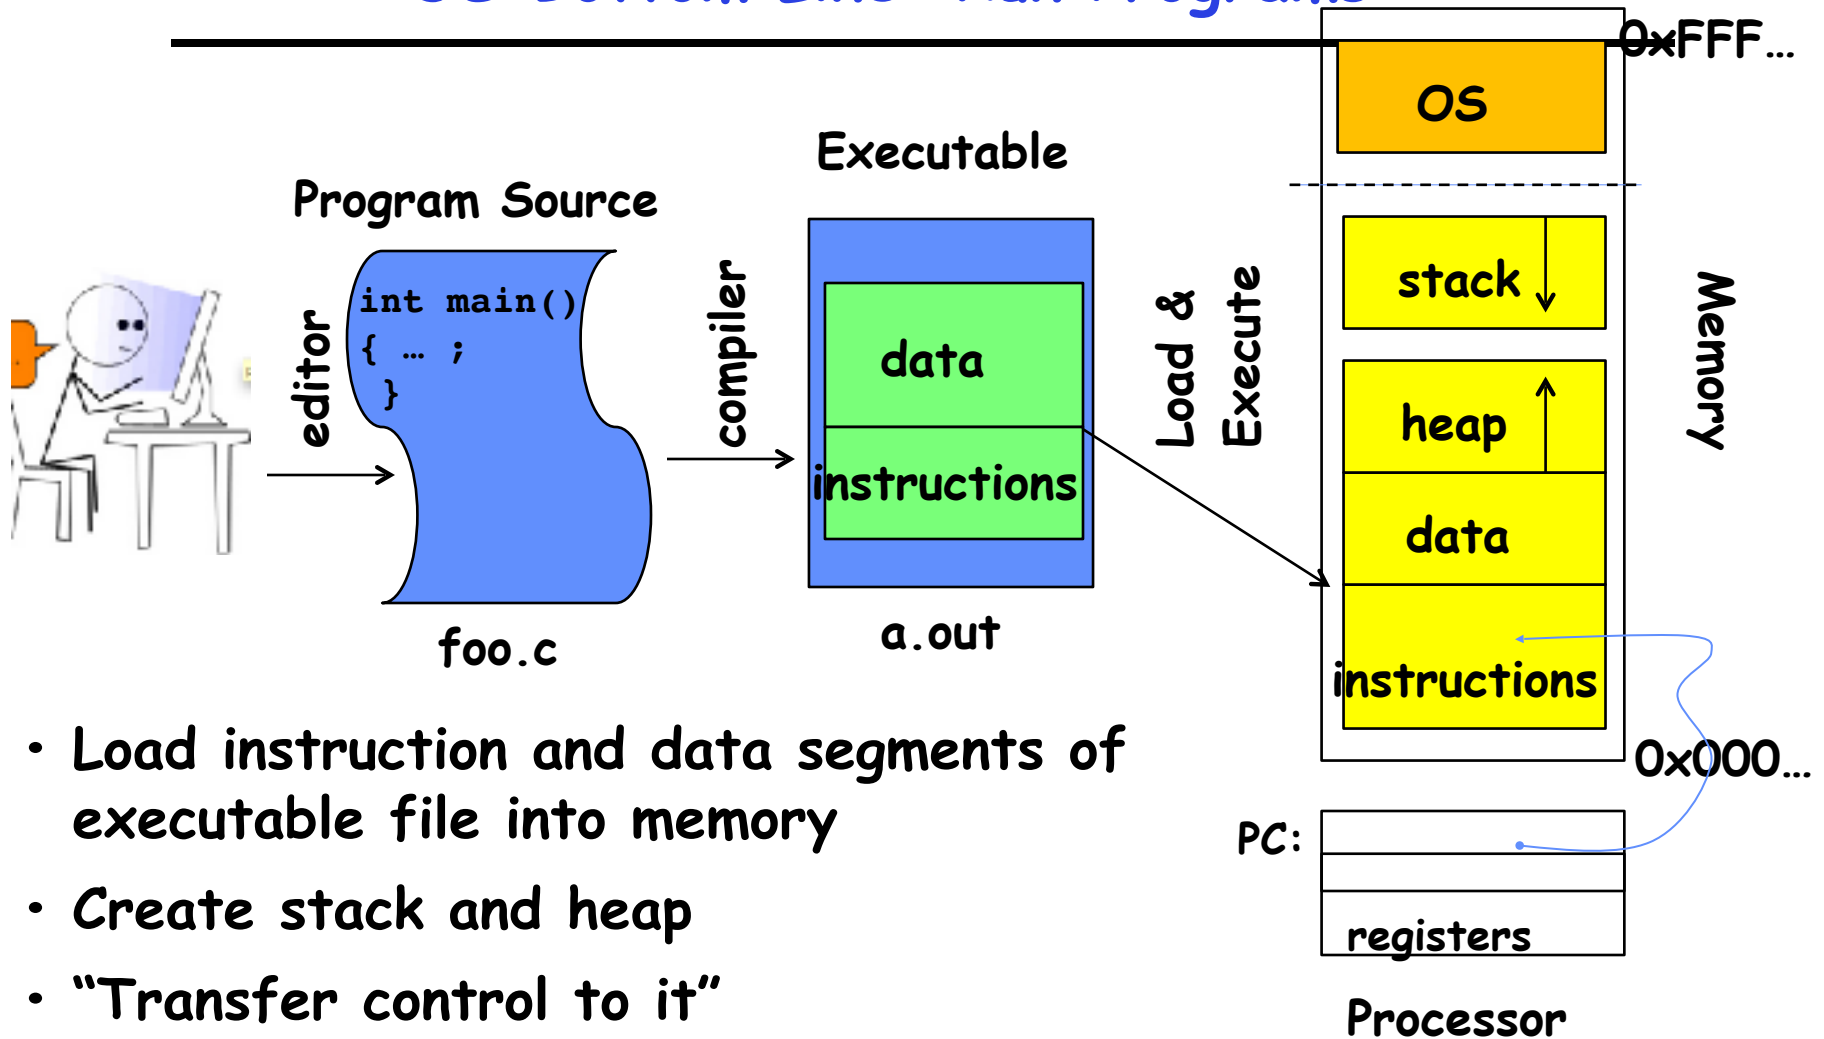
- Because of the cost of developing an OS from scratch, most modern OSes have a long lineage:
- Multics → AT&T Unix → BSD Unix → Ultrix, SunOS, NetBSD,...
- Mach (micro-kernel) + BSD → NextStep → XNU → Apple OSX, iPhone iOS
- Linux → Android OS
- CP/M → QDOS → MS-DOS → Windows 3.1 → NT → 95 → 98 → 2000 → XP → Vista → 7 → 8 → phone → ...
- Linux → RedHat, Ubuntu, Fedora, Debian, Suse,...

# Today: Four fundamental OS concepts

---

- **Thread**
  - Single unique execution context
  - Program Counter, Registers, Execution Flags, Stack
- **Address Space w/ Translation**
  - Programs execute in an address space that is distinct from the memory space of the physical machine
- **Process**
  - An instance of an executing program is a process consisting of an address space and one or more threads of control
- **Dual Mode operation/Protection**
  - Only the "system" has the ability to access certain resources
  - The OS and the hardware are protected from user programs and user programs are isolated from one another by controlling the translation from program virtual addresses to machine physical addresses

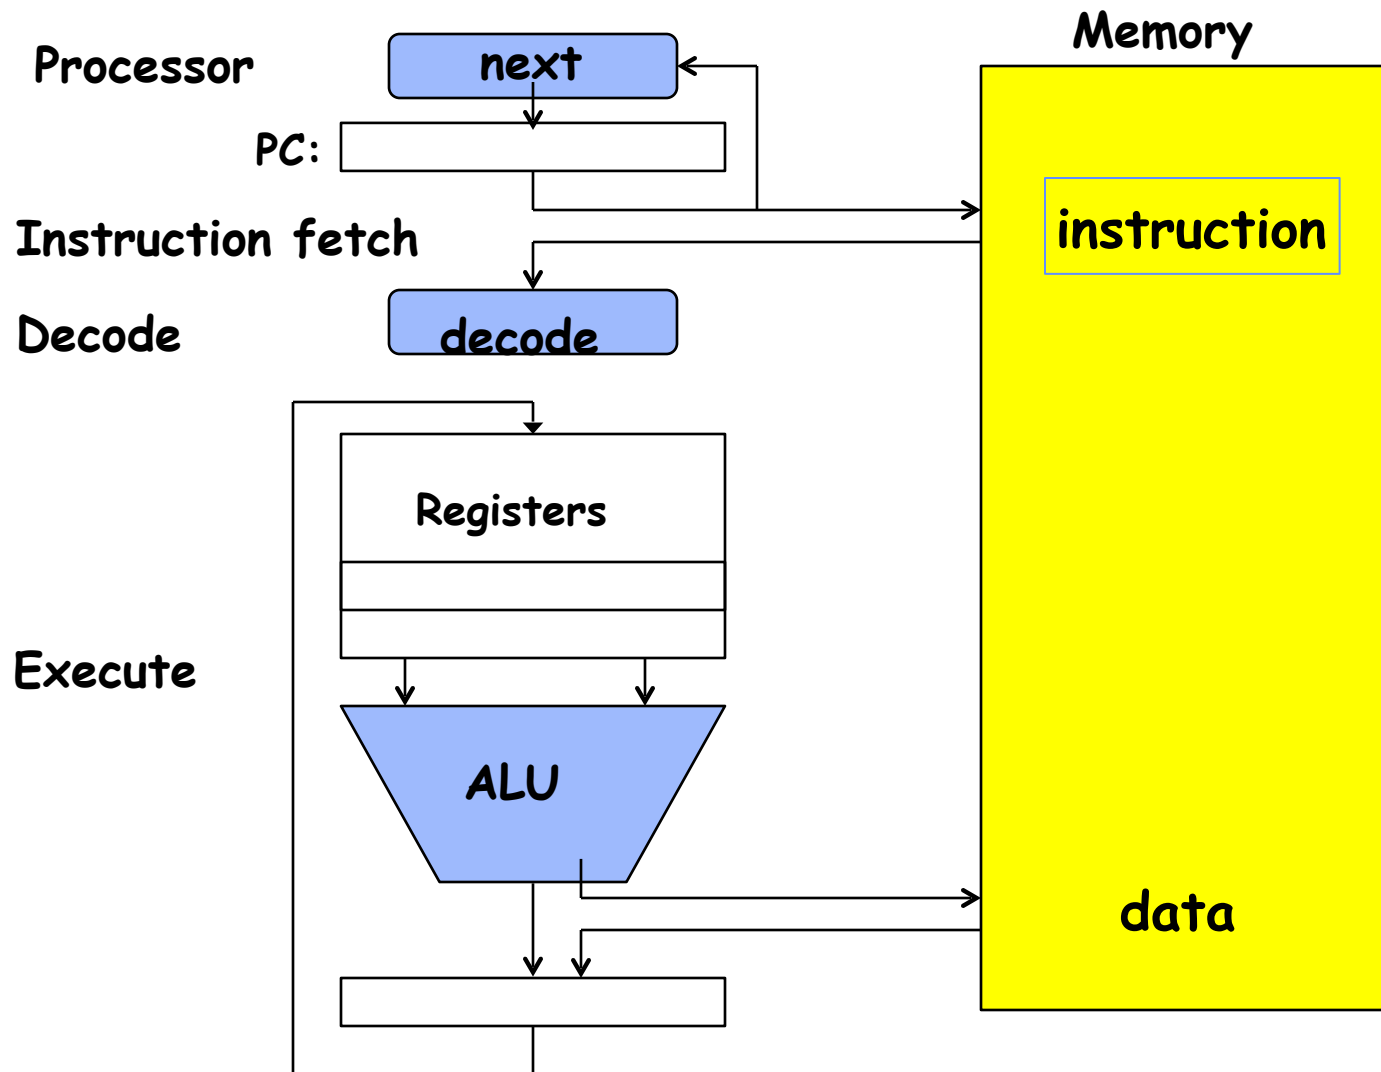
# OS Bottom Line: Run Programs



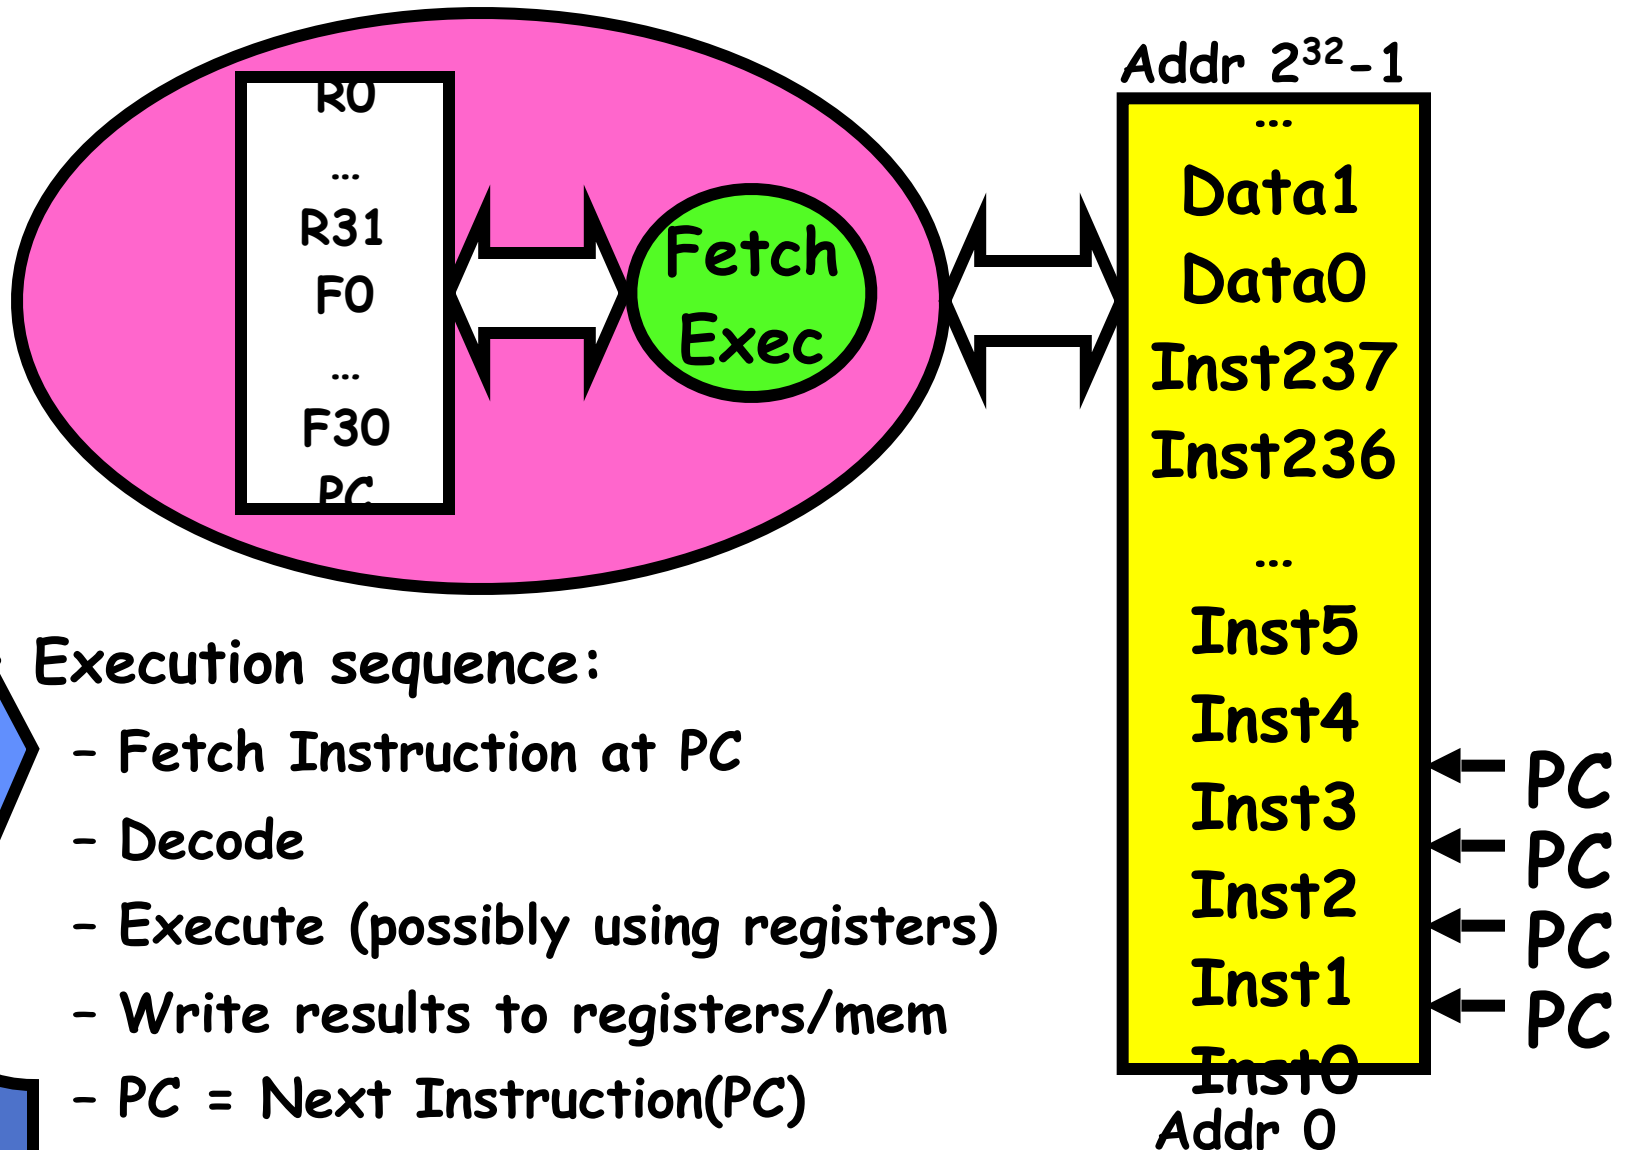
- Load instruction and data segments of executable file into memory
- Create stack and heap
- "Transfer control to it"
- Provide services to it
- While protecting OS and it

# Today we need one key concept

## The instruction cycle




# Recall: What happens during program execution?



## Execution sequence:

- Fetch Instruction at PC
- Decode
- Execute (possibly using registers)
- Write results to registers/mem
- PC = Next Instruction(PC)
- Repeat




# First OS Concept: Thread of Control

---

- **Thread: Single unique execution context**
  - Program Counter, Registers, Execution Flags, Stack
- A thread is executing on a processor when it is resident in the processor registers.
- PC register holds the address of executing instruction in the thread.
- Certain registers hold the context of thread
  - Stack pointer holds the address of the top of stack
    - » Other conventions: Frame Pointer, Heap Pointer, Data
  - May be defined by the instruction set architecture or by compiler conventions
- Registers hold the root state of the thread.
  - The rest is "in memory"

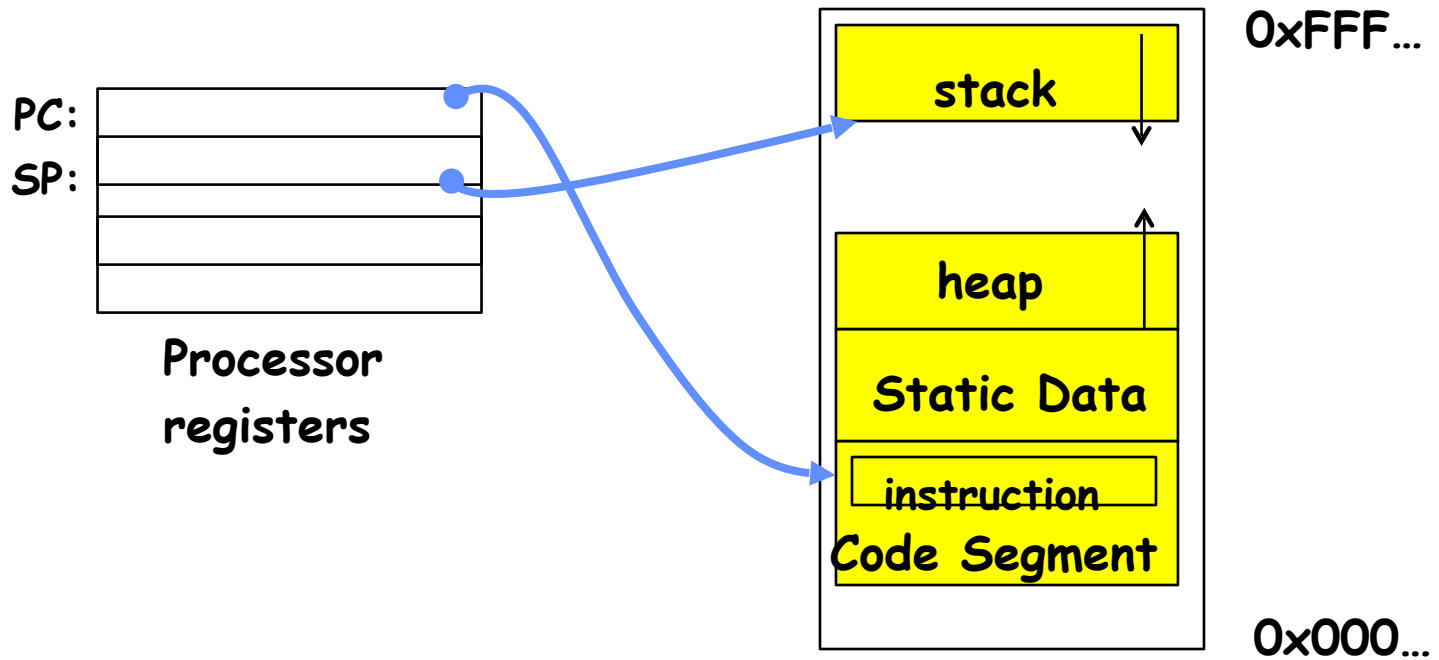
## Second OS Concept: Program's Address Space

- **Address space**  $\Rightarrow$  the set of accessible addresses + state associated with them:
  - For a 32-bit processor there are  $2^{32} = 4$  billion addresses
- What happens when you read or write to an address?
  - Perhaps Nothing
  - Perhaps acts like regular memory
  - Perhaps ignores writes
  - Perhaps causes I/O operation
    - » (Memory-mapped I/O)
  - Perhaps causes exception (fault)



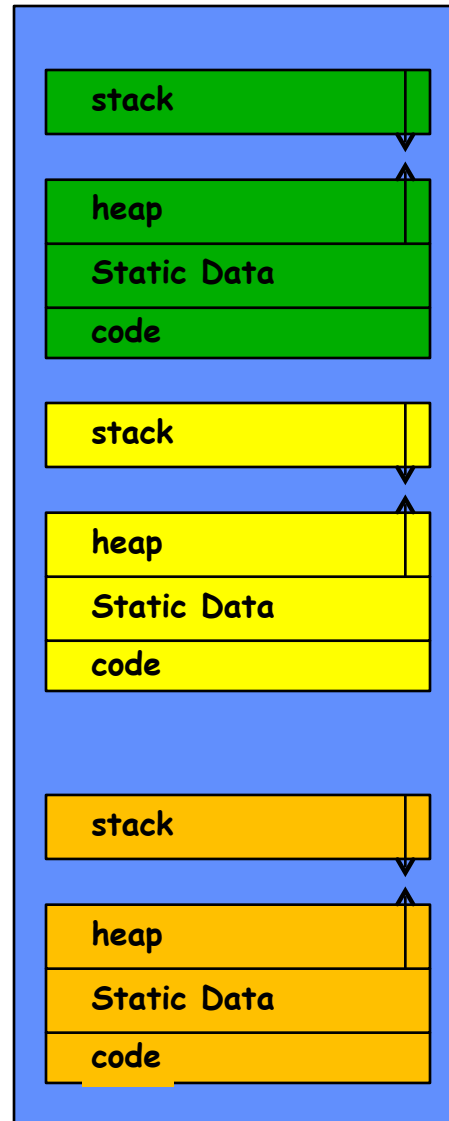
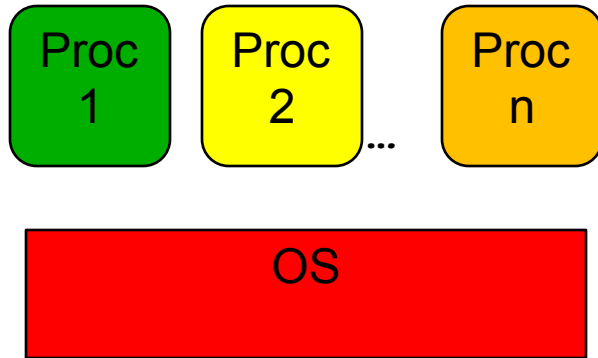


# Address Space: In a Picture



- What's in the code segment? Data?
- What's in the stack segment?
  - How is it allocated? How big is it?
- What's in the heap segment?
  - How is it allocated? How big?

# Multiprogramming - Multiple Threads of Control



# Administrivia: Getting started

---

- Start homework 0 immediately ⇒ **Due Mehr 5th**
  - Github account
  - Vagrant virtualbox - VM environment for the course
    - » Consistent, managed environment on your machine
  - Get familiar with all the tools
- Office hours
  - Monday 1630 to 1700, email me for other times
- TA Class
  - Saturday 12 to 13, or Monday 12 to 13?
- HW/GHW schedule
  - Do consider murphy's law
  - things break when you are not expecting it, so plan for it
- Group sign up form out next week (after drop deadline)
  - Find group members ASAP
  - 4 people in a group!

# CS 162 Collaboration Policy

---

Explaining a concept to someone in another group

 Discussing algorithms/testing strategies with other groups

Helping debug someone else's code (in another group)

Searching online for generic algorithms (e.g., hash table)


 Sharing code or test cases with another group

Copying OR reading another group's code or test cases

Copying OR reading online code or test cases from from prior years

We compare all project submissions against prior year submissions and online solutions and will take actions (described on the course overview page) against offenders

# How can we give the illusion of multiple processors?



- Assume a single processor. How do we provide the illusion of multiple processors?
  - Multiplex in time!
- Each virtual "CPU" needs a structure to hold:
  - Program Counter (PC), Stack Pointer (SP)
  - Registers (Integer, Floating point, others...?)
- How switch from one virtual CPU to the next?
  - Save PC, SP, and registers in current state block
  - Load PC, SP, and registers from new state block
- What triggers switch?
  - Timer, voluntary yield, I/O, other things

# The Basic Problem of Concurrency

---

- The basic problem of concurrency involves resources:
  - Hardware: single CPU, single DRAM, single I/O devices
  - Multiprogramming API: processes think they have exclusive access to shared resources
- OS has to coordinate all activity
  - Multiple processes, I/O interrupts, ...
  - How can it keep all these things straight?
- Basic Idea: Use Virtual Machine abstraction
  - Simple machine abstraction for processes
  - Multiplex these abstract machines
- Dijkstra did this for the "THE system"
  - Few thousand lines vs 1 million lines in OS 360 (1K bugs)

## Properties of this simple multiprogramming technique

- All virtual CPUs share same non-CPU resources
  - I/O devices the same
  - Memory the same
- Consequence of sharing:
  - Each thread can access the data of every other thread (good for sharing, bad for protection)
  - Threads can share instructions (good for sharing, bad for protection)
  - Can threads overwrite OS functions?
- This (unprotected) model is common in:
  - Embedded applications
  - Windows 3.1/Early Macintosh (switch only with yield)
  - Windows 95—ME (switch with both yield and timer)

## Third OS Concept: Process

---

- **Process: execution environment with Restricted Rights**
  - **Address Space with One or More Threads**
  - Owns memory (address space)
  - Owns file descriptors, file system context, ...
  - Encapsulate one or more threads sharing process resources
- **Why processes?**
  - **Protected from each other!**
  - **OS Protected from them**
  - Processes provides memory protection
  - Threads more efficient than processes (later)
- **Fundamental tradeoff between protection and efficiency**
  - Communication easier within a process
  - Communication harder between processes
- **Application instance consists of one or more processes**



# Protection

---


- **Operating System must protect itself from user programs**
  - **Reliability:** compromising the operating system generally causes it to crash
  - **Security:** limit the scope of what processes can do
  - **Privacy:** limit each process to the data it is permitted to access
  - **Fairness:** each should be limited to its appropriate share of system resources (CPU time, memory, I/O, etc)
- **It must protect User programs from one another**
- **Primary Mechanism:** limit the translation from program address space to physical memory space
  - Can only touch what is mapped into process address space
- **Additional Mechanisms:**
  - Privileged instructions, in/out instructions, special registers
  - syscall processing, subsystem implementation
    - » (e.g., file access rights, etc)

## Fourth OS Concept: Dual Mode Operation

---


- **Hardware** provides at least two modes:
  - “Kernel” mode (or “supervisor” or “protected”)
  - “User” mode: Normal programs executed
- What is needed in the hardware to support “dual mode” operation?
  - a bit of state (user/system mode bit)
  - Certain operations / actions only permitted in system/kernel mode
    - » In user mode they fail or trap
  - User-→Kernel transition sets system mode AND saves the user PC
    - » Operating system code carefully puts aside user state then performs the necessary operations
  - Kernel-→User transition clears system mode AND restores appropriate user PC
    - » return-from-interrupt

# For example: UNIX System Structure

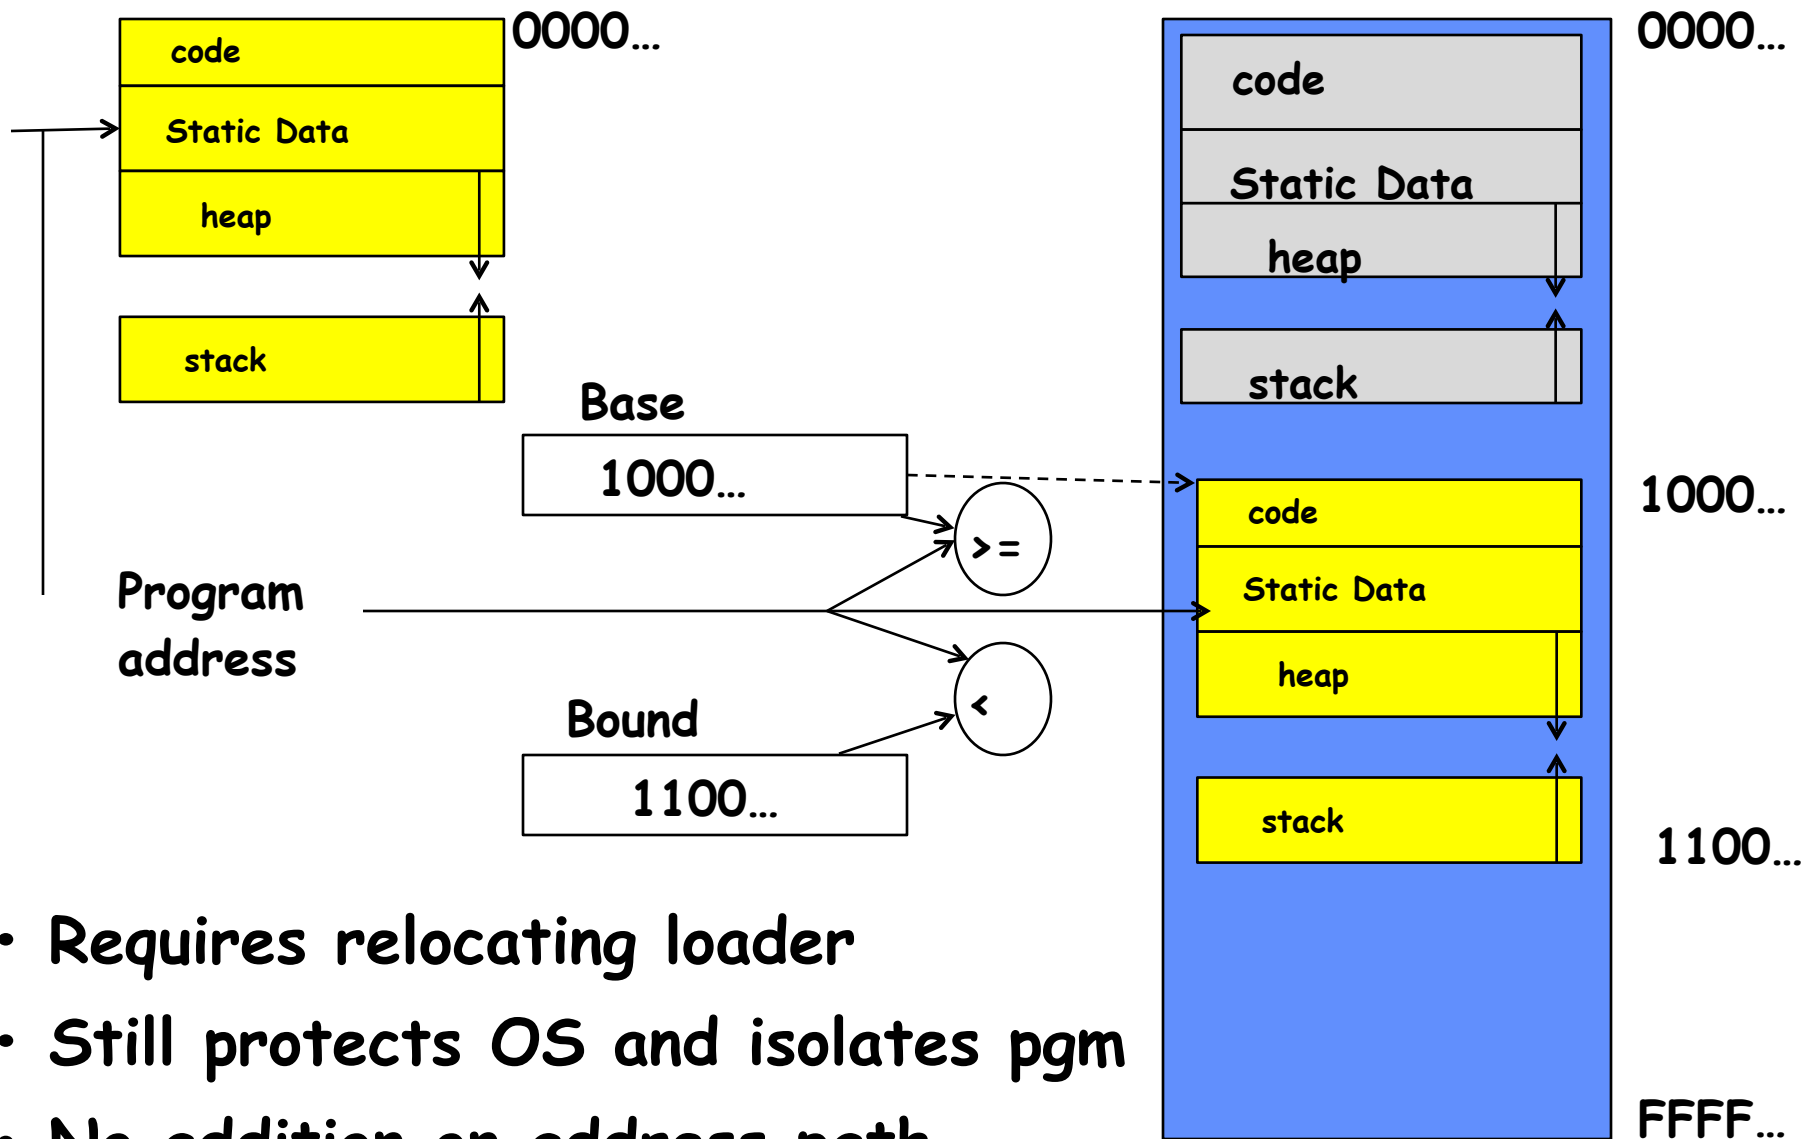


# User/Kernel(Privileged) Mode

---



# Simple Protection: Base and Bound (B&B)

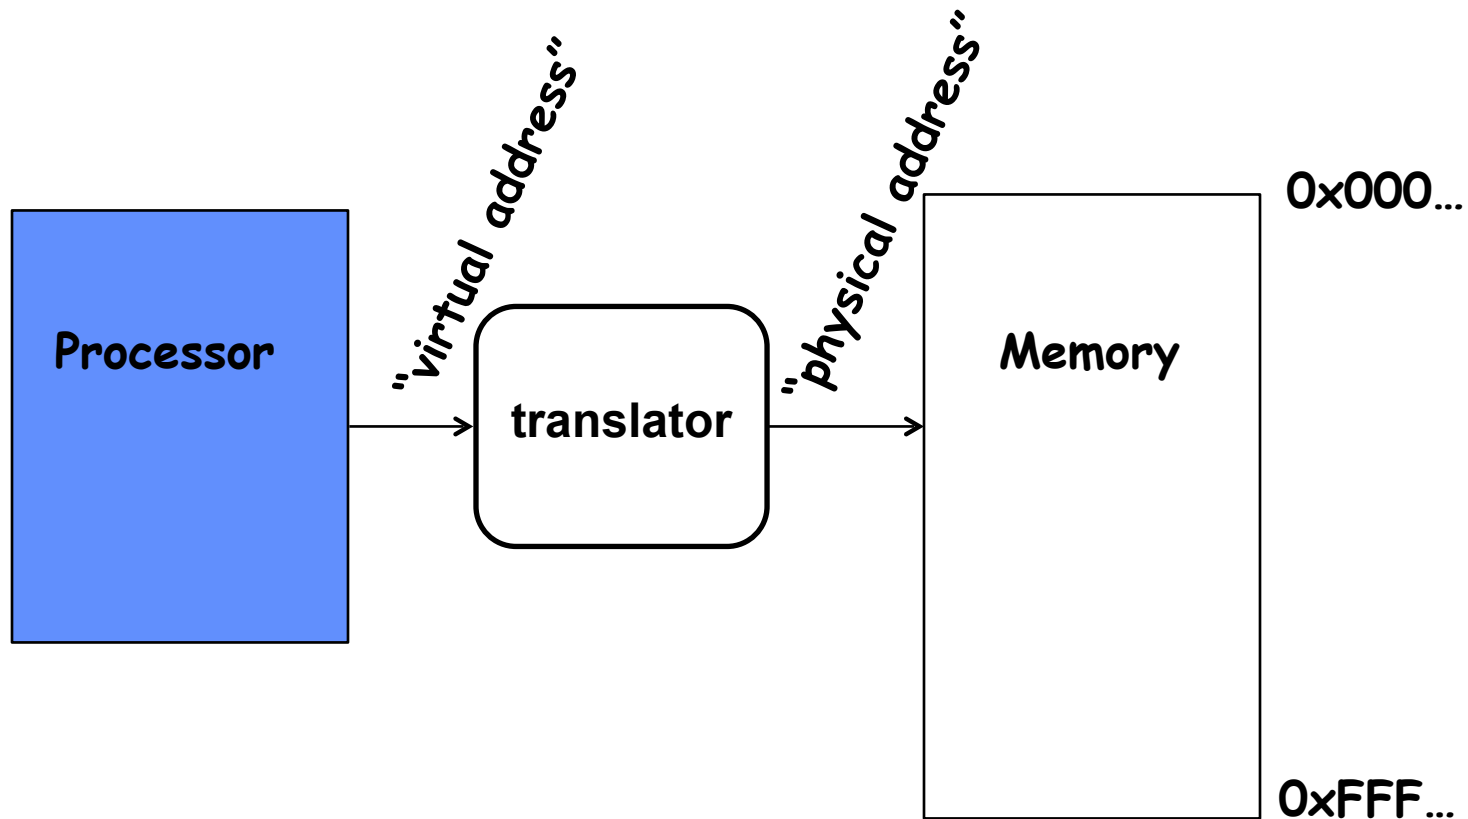


- Requires relocating loader
- Still protects OS and isolates pgm
- No addition on address path

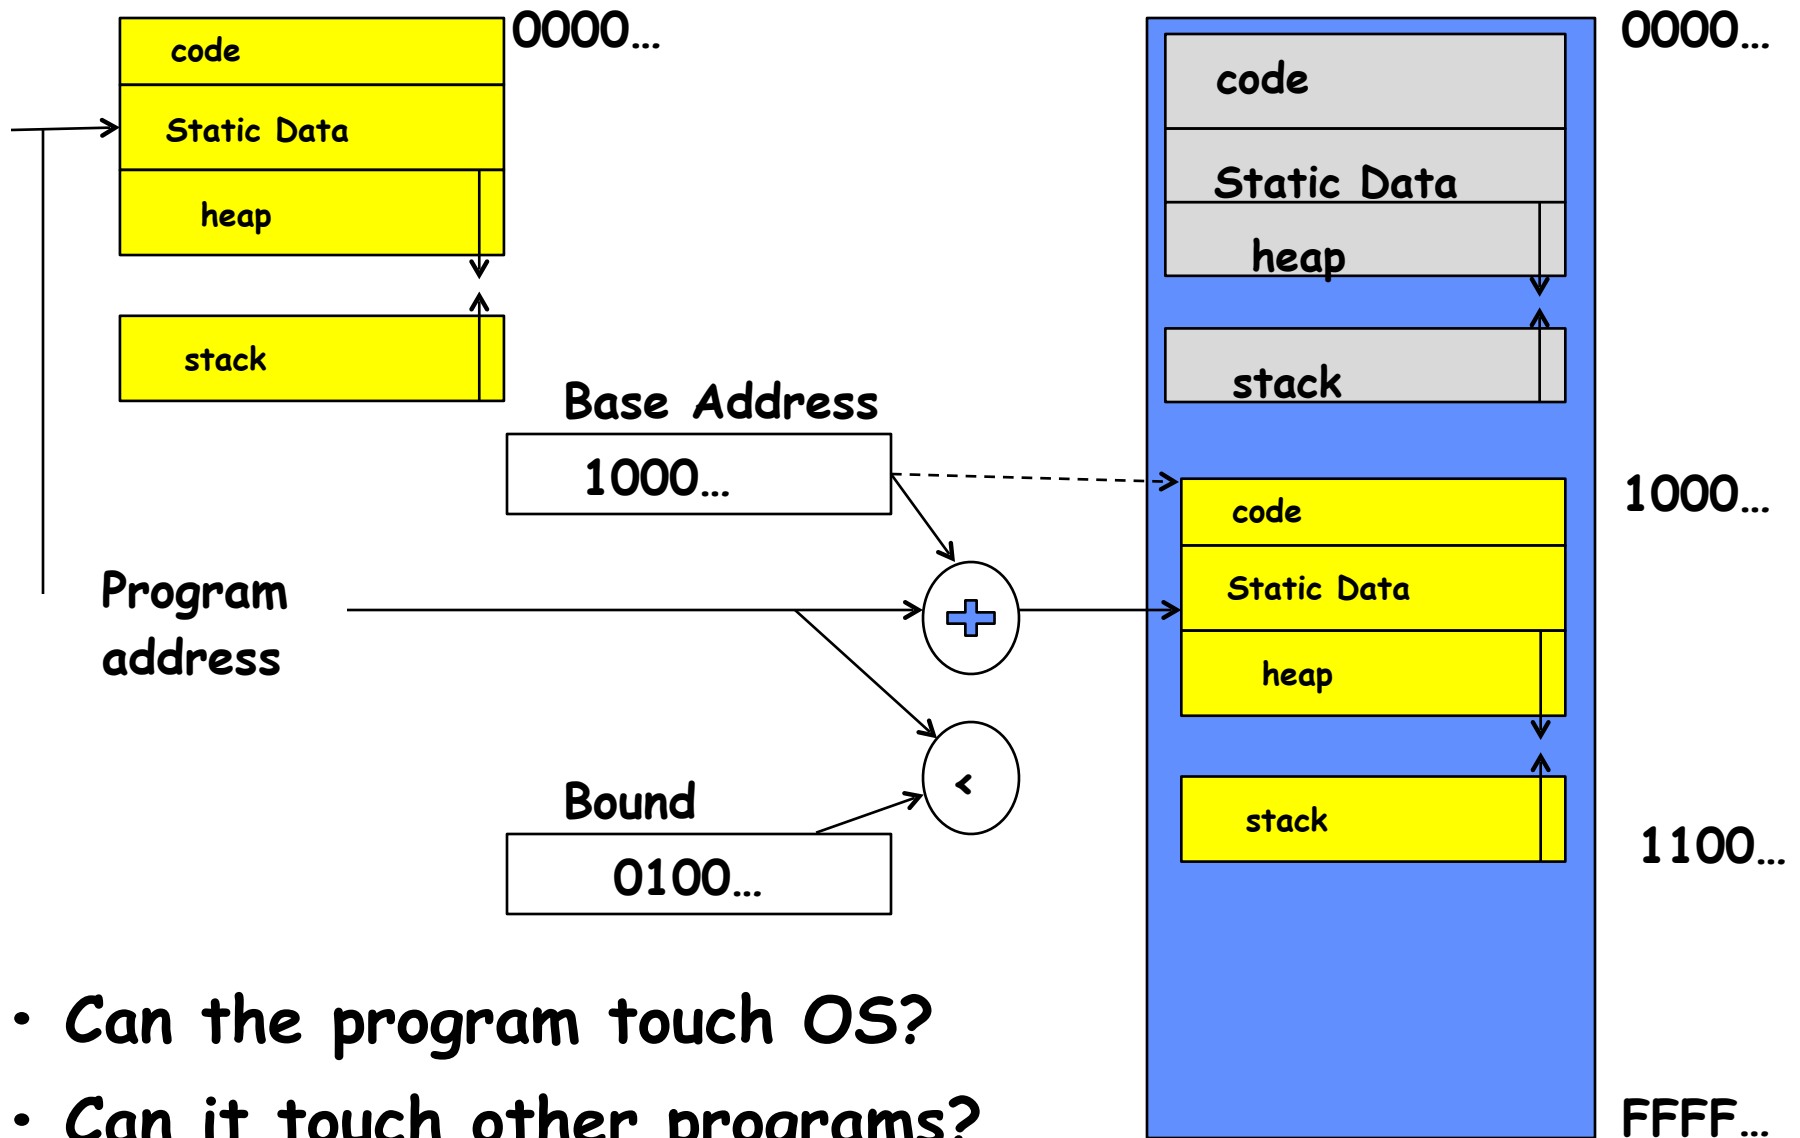
## Another idea: Address Space Translation

---

- Program operates in an address space that is distinct from the physical memory space of the machine

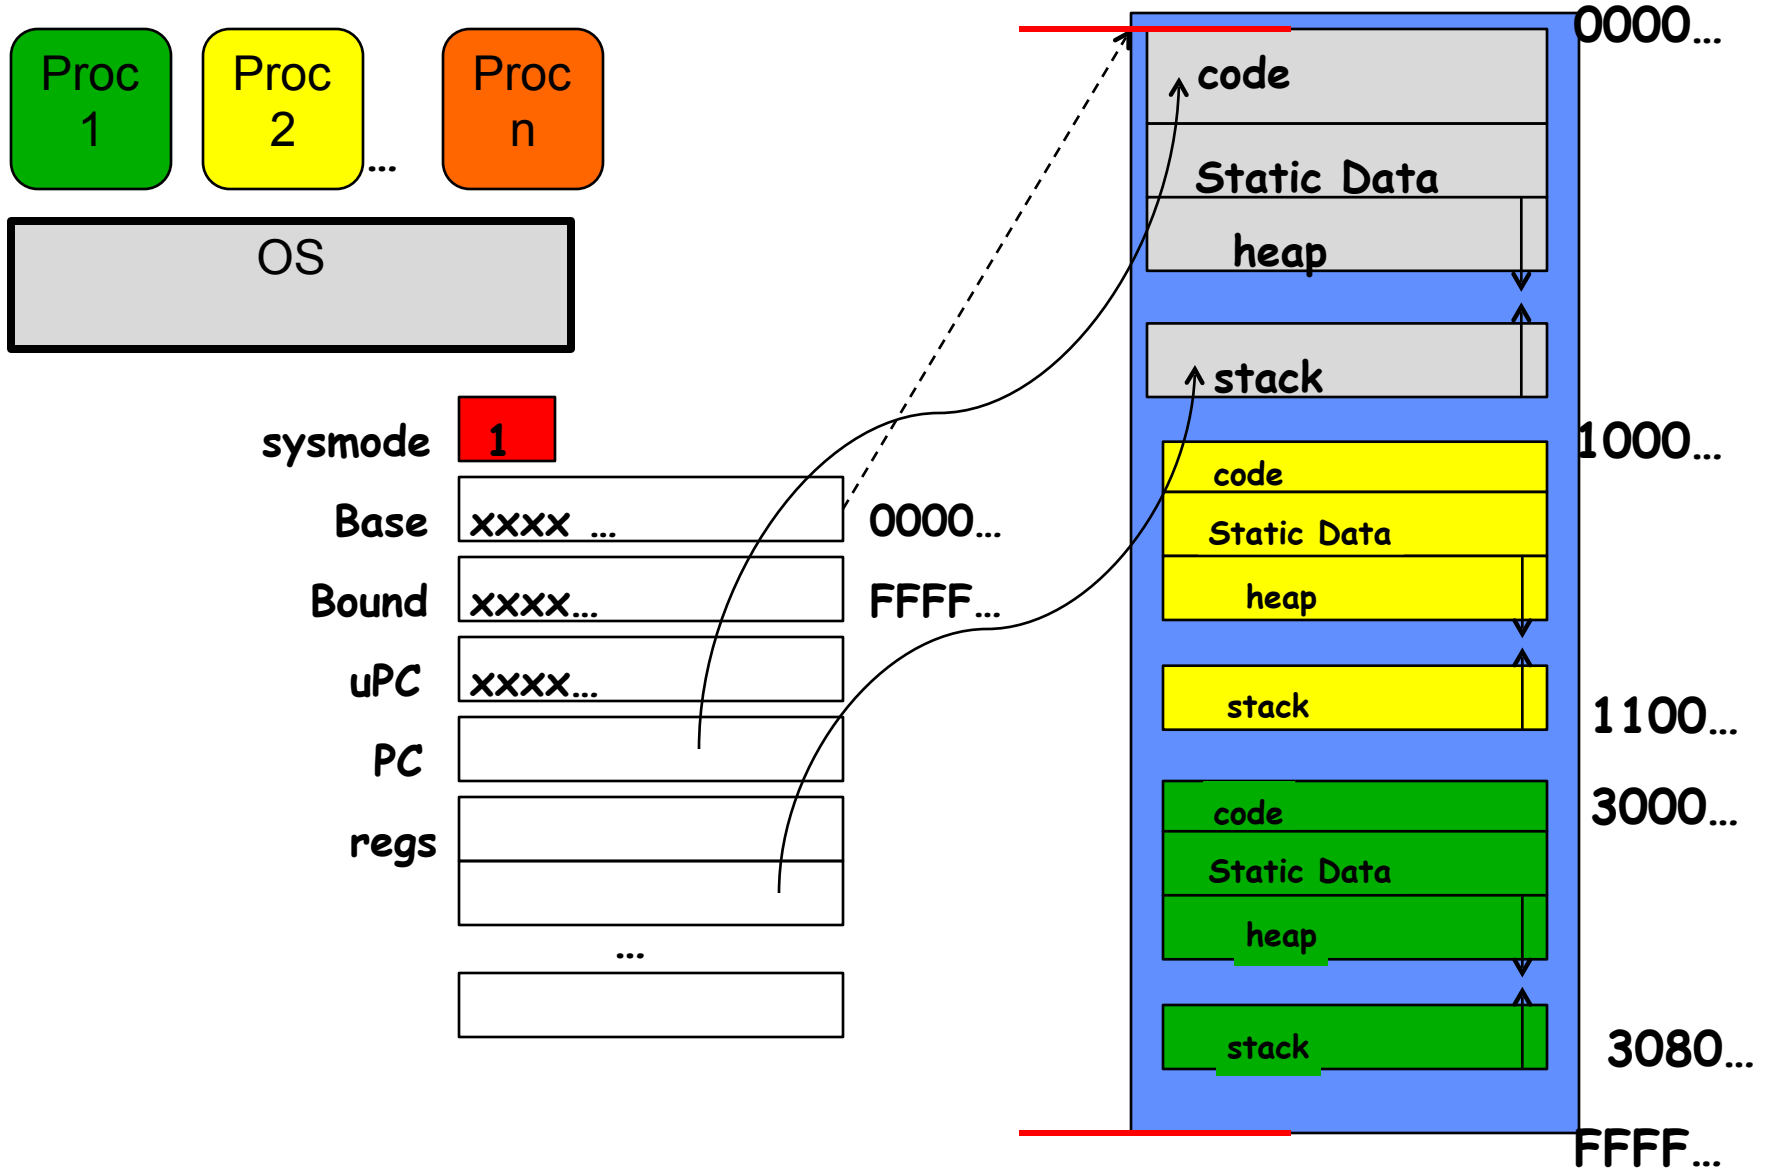


# A simple address translation with Base and Bound



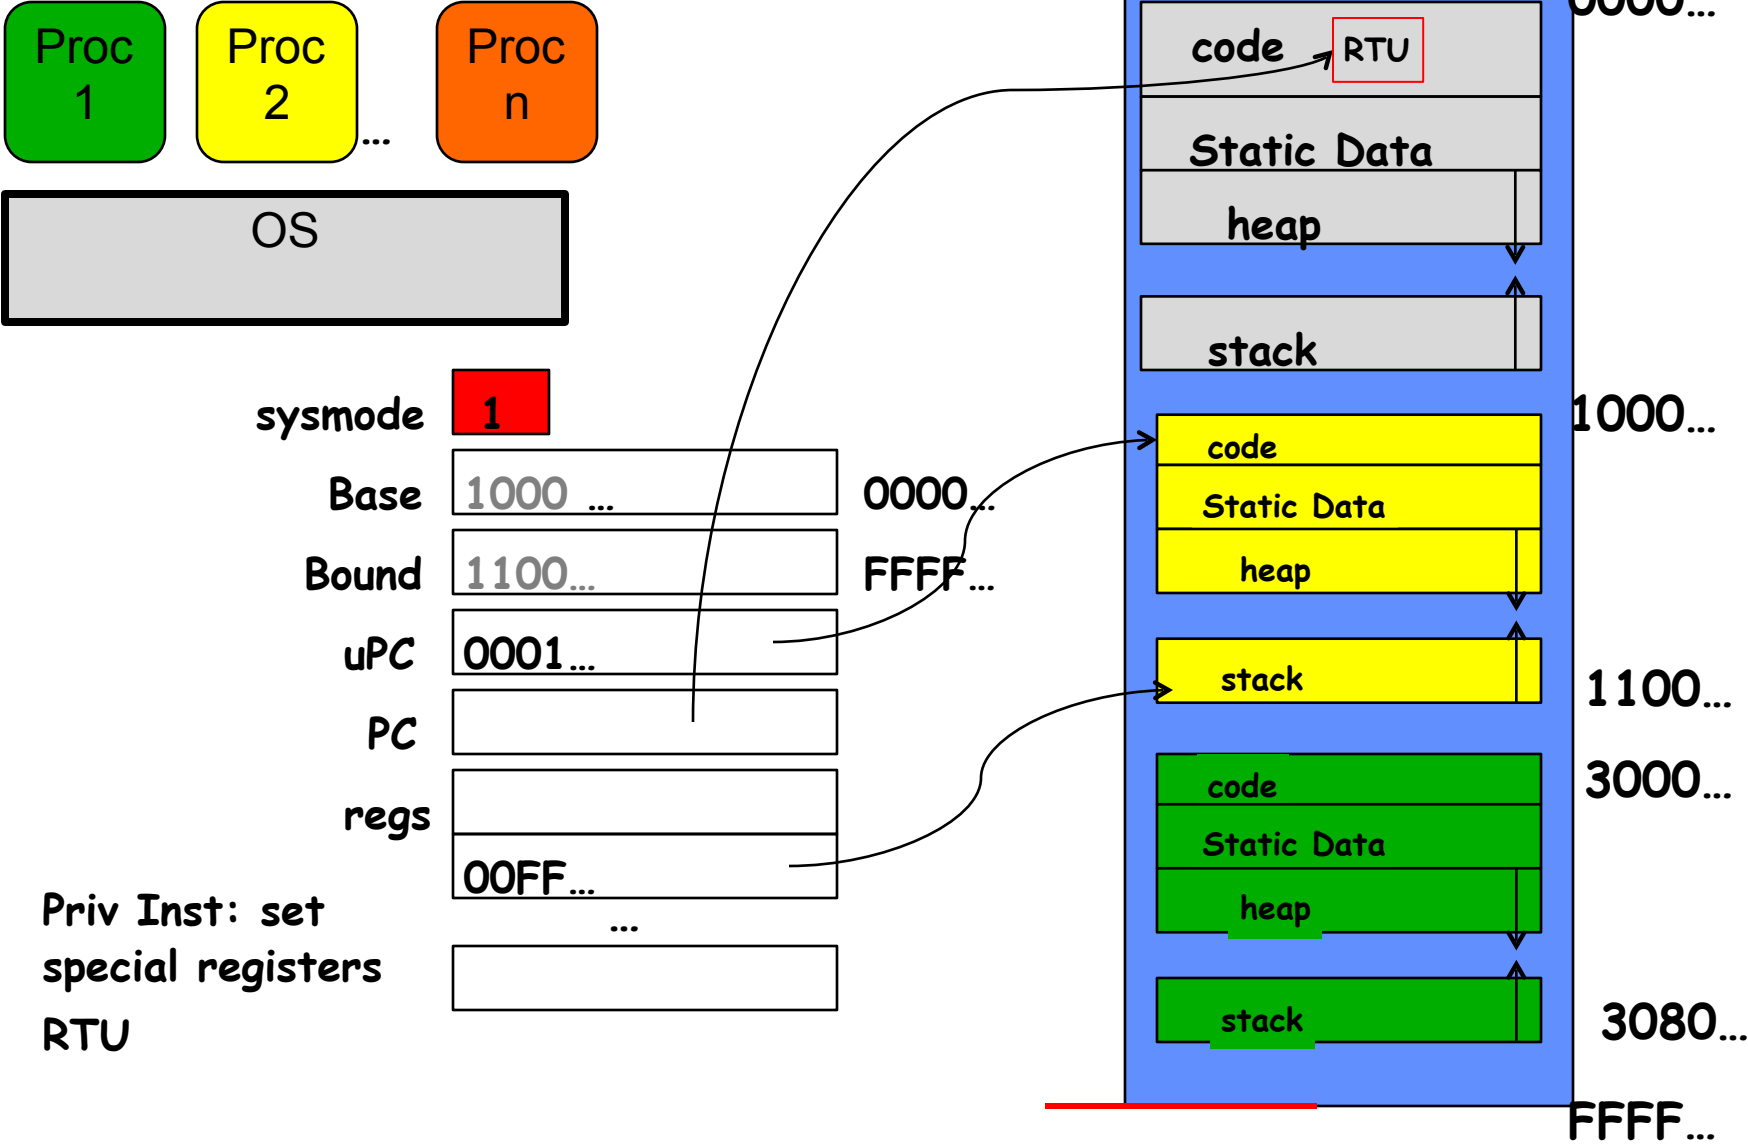
- Can the program touch OS?
- Can it touch other programs?

# Tying it together: Simple B&B: OS loads process



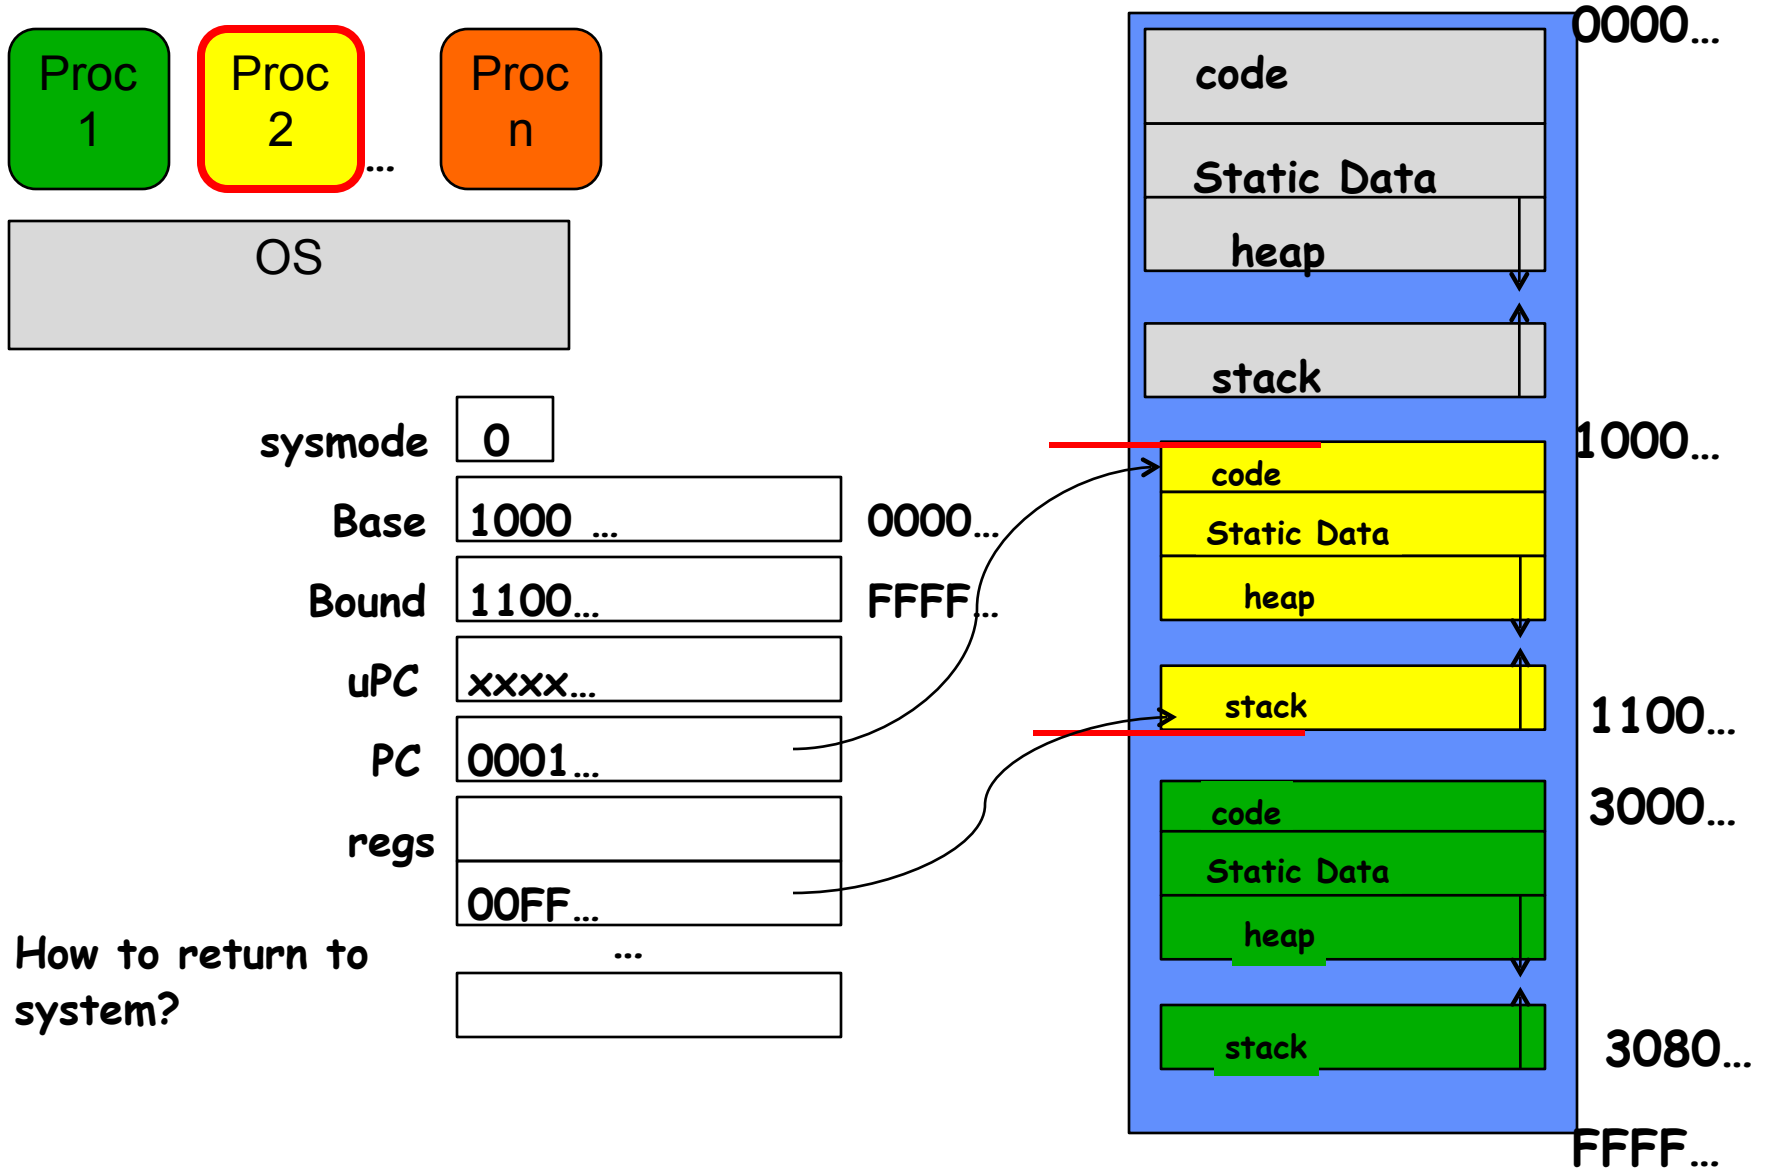


# Simple B&B: OS gets ready to switch



- Priv Inst: set special registers
- RTU

# Simple B&B: "Return" to User



# 3 types of Mode Transfer

---

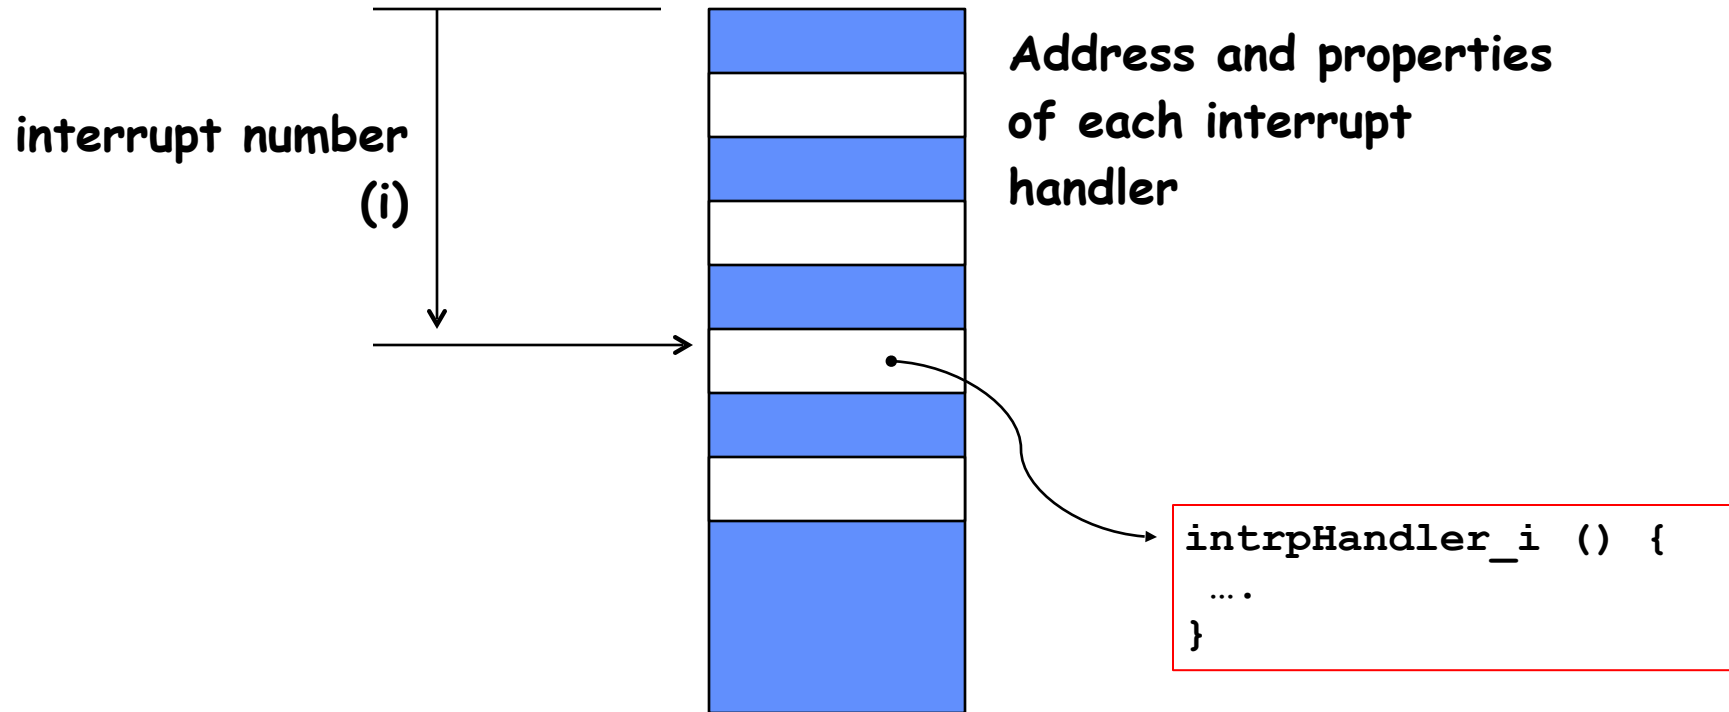
- **Syscall**
  - Process requests a system service, e.g., exit
  - Like a function call, but “outside” the process
  - Does not have the address of the system function to call
  - Like a Remote Procedure Call (RPC) - for later
  - Marshall the syscall id and args in registers and exec syscall
- **Interrupt**
  - External asynchronous event triggers context switch
  - eg. Timer, I/O device
  - Independent of user process
- **Trap or Exception**
  - Internal synchronous event in process triggers context switch
  - e.g., Protection violation (segmentation fault), Divide by zero, ...
- **All 3 are an UNPROGRAMMED CONTROL TRANSFER**
  - Where does it go?

---

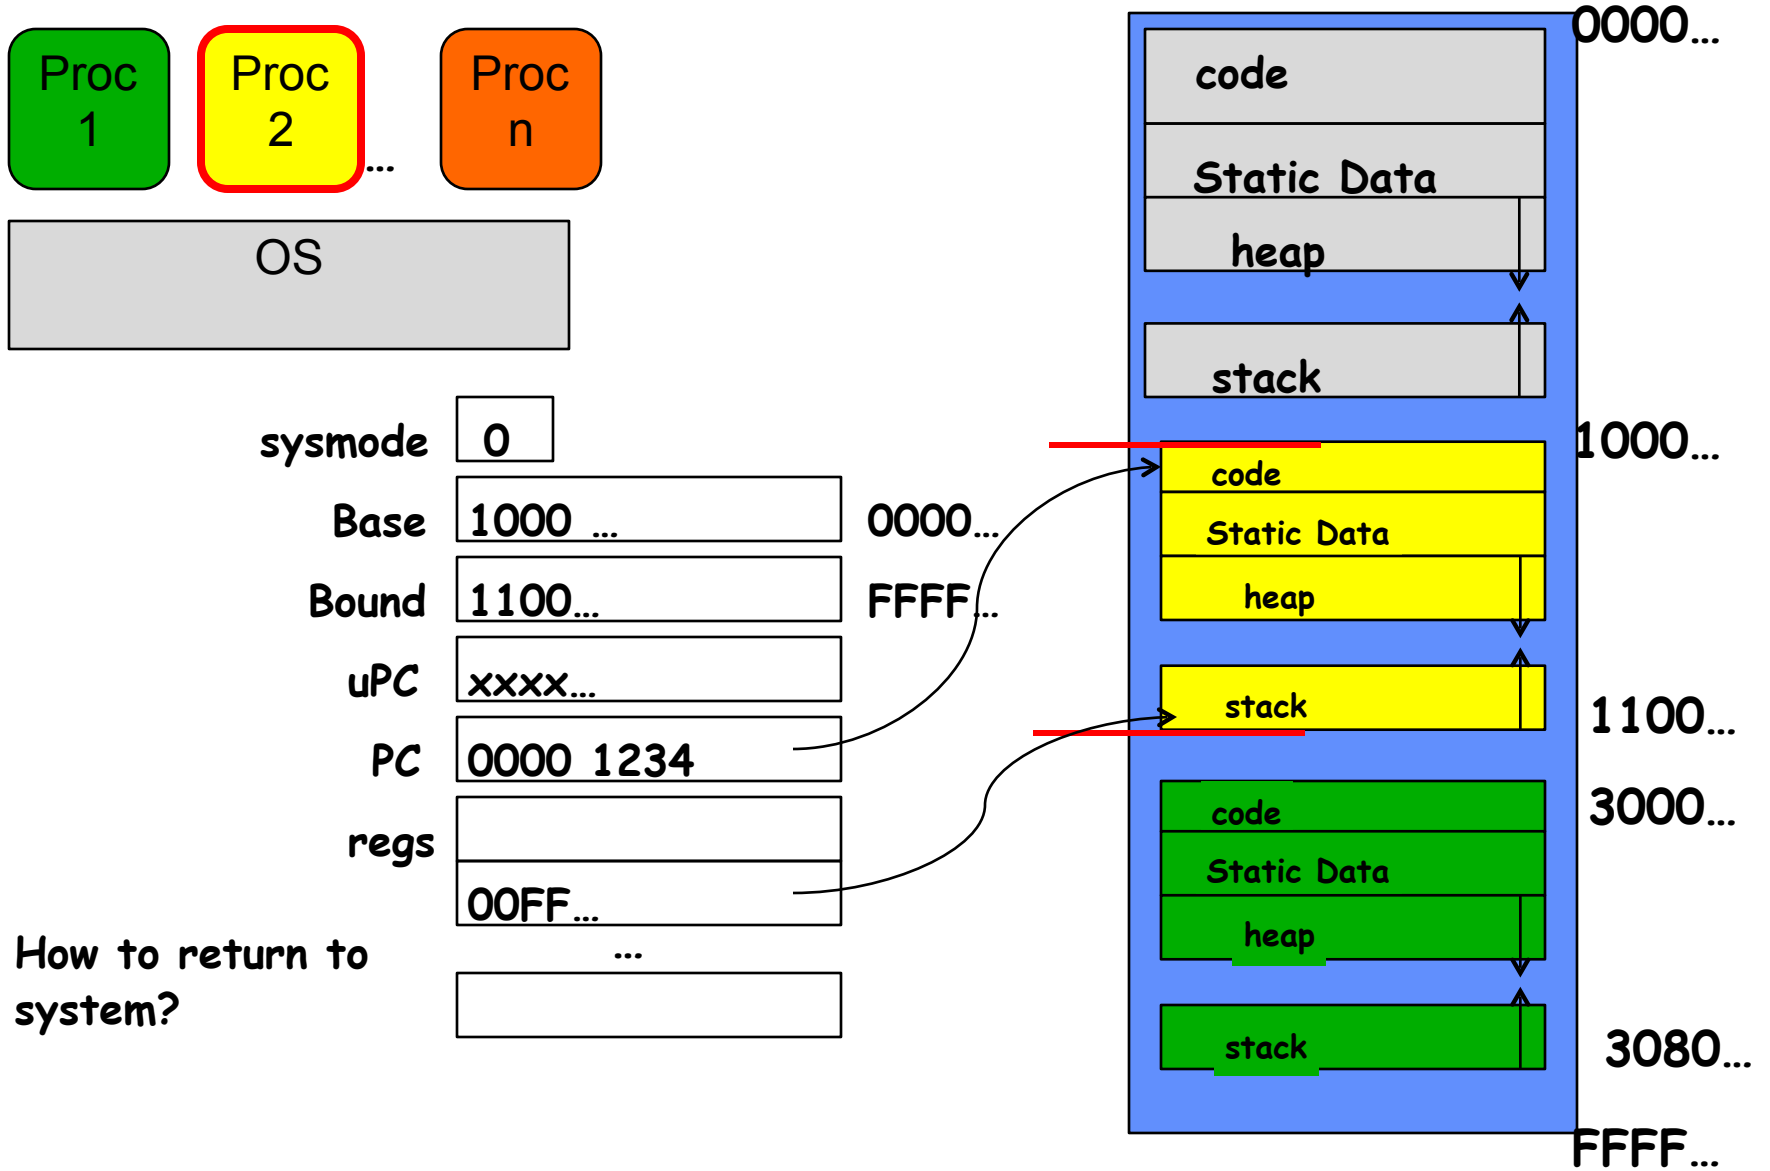
How do we get the system target address of the “unprogrammed control transfer?”

# Interrupt Vector

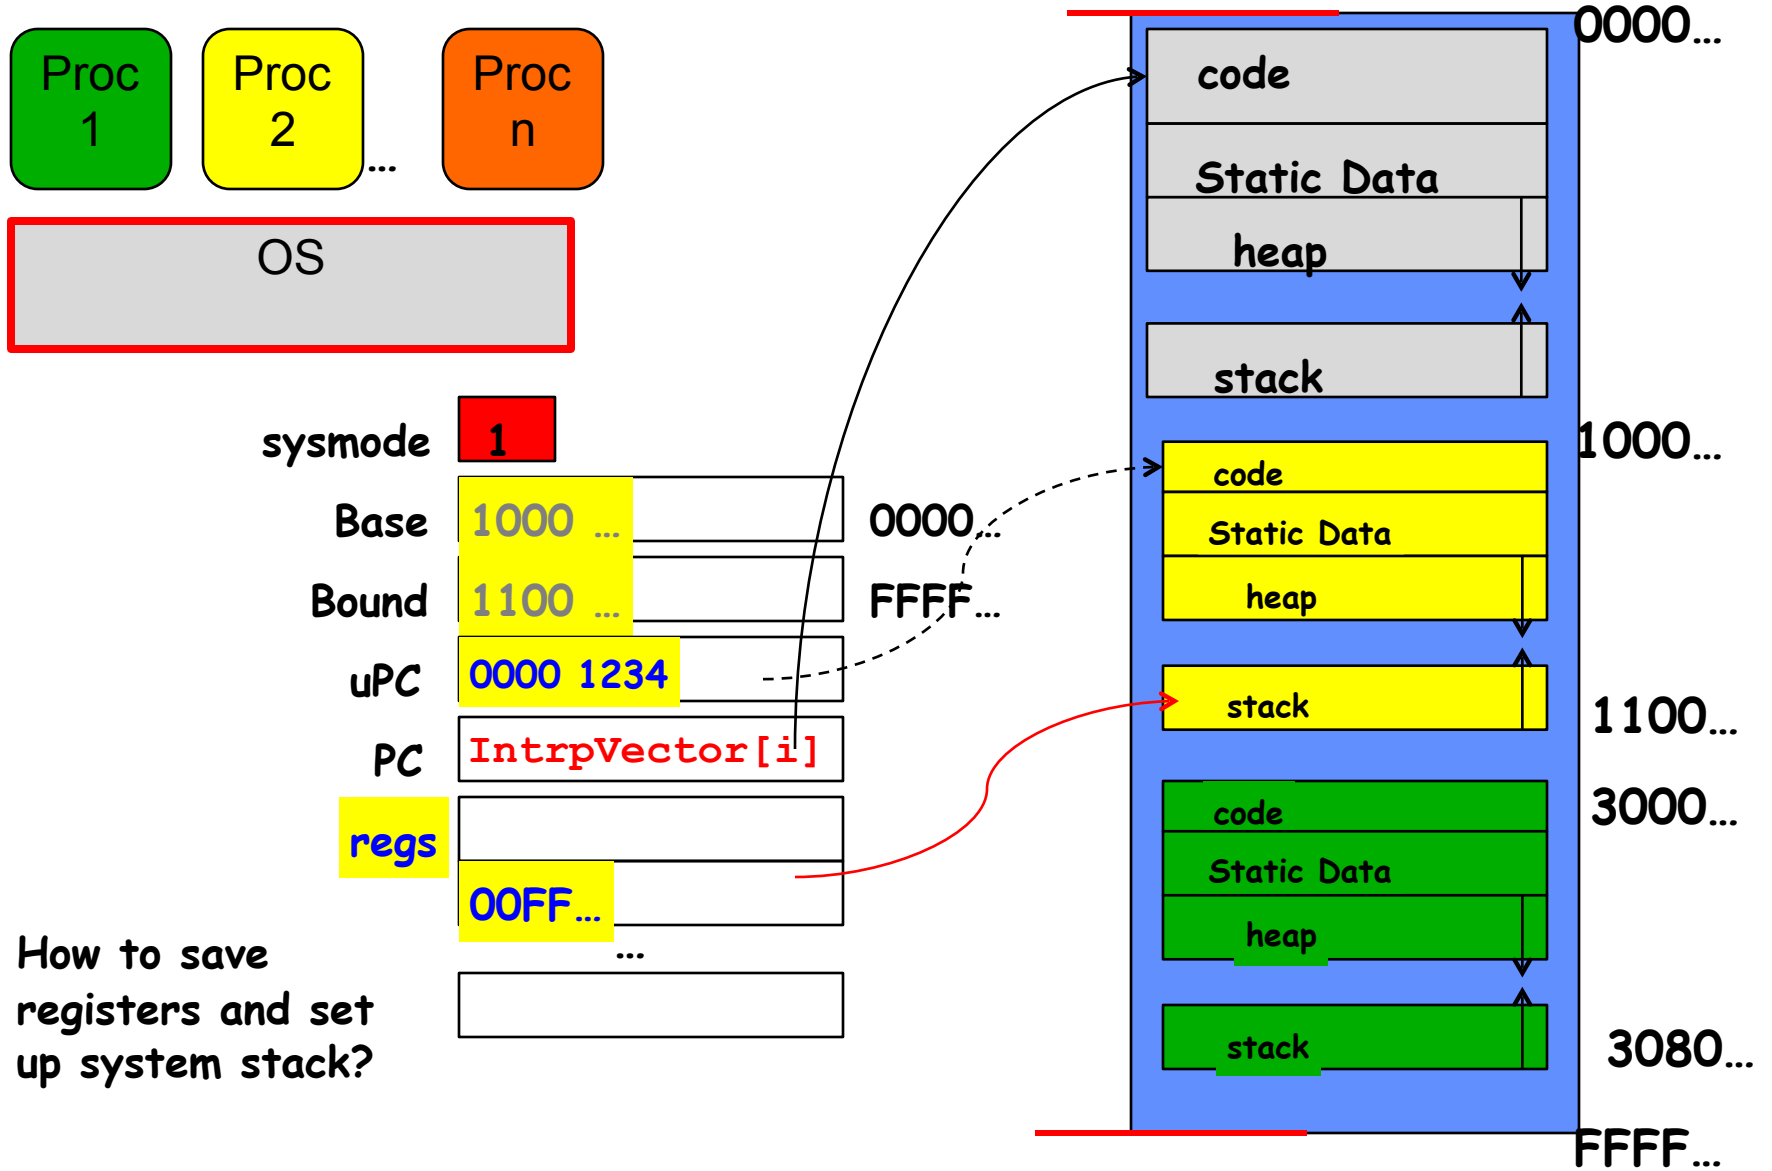
---



# Simple B&B: User => Kernel

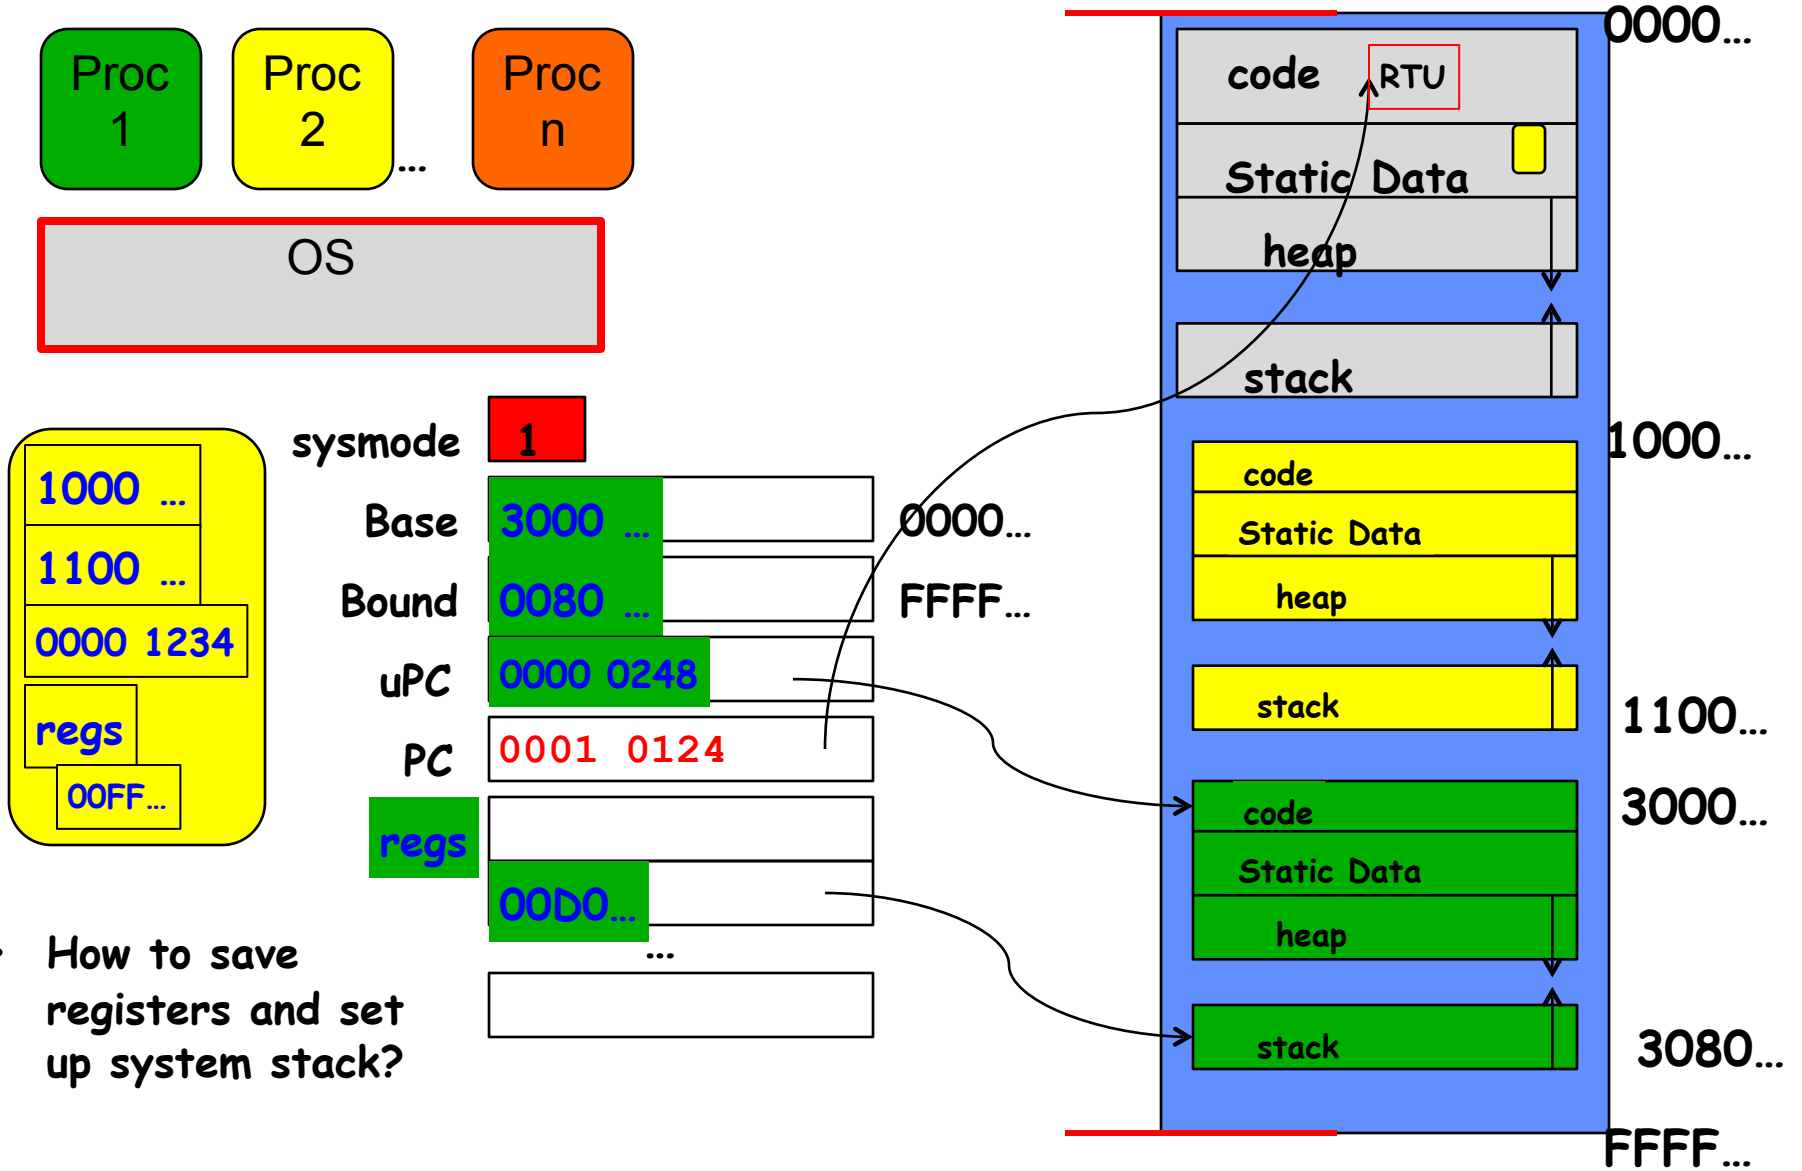


# Simple B&B: Interrupt



- How to save registers and set up system stack?

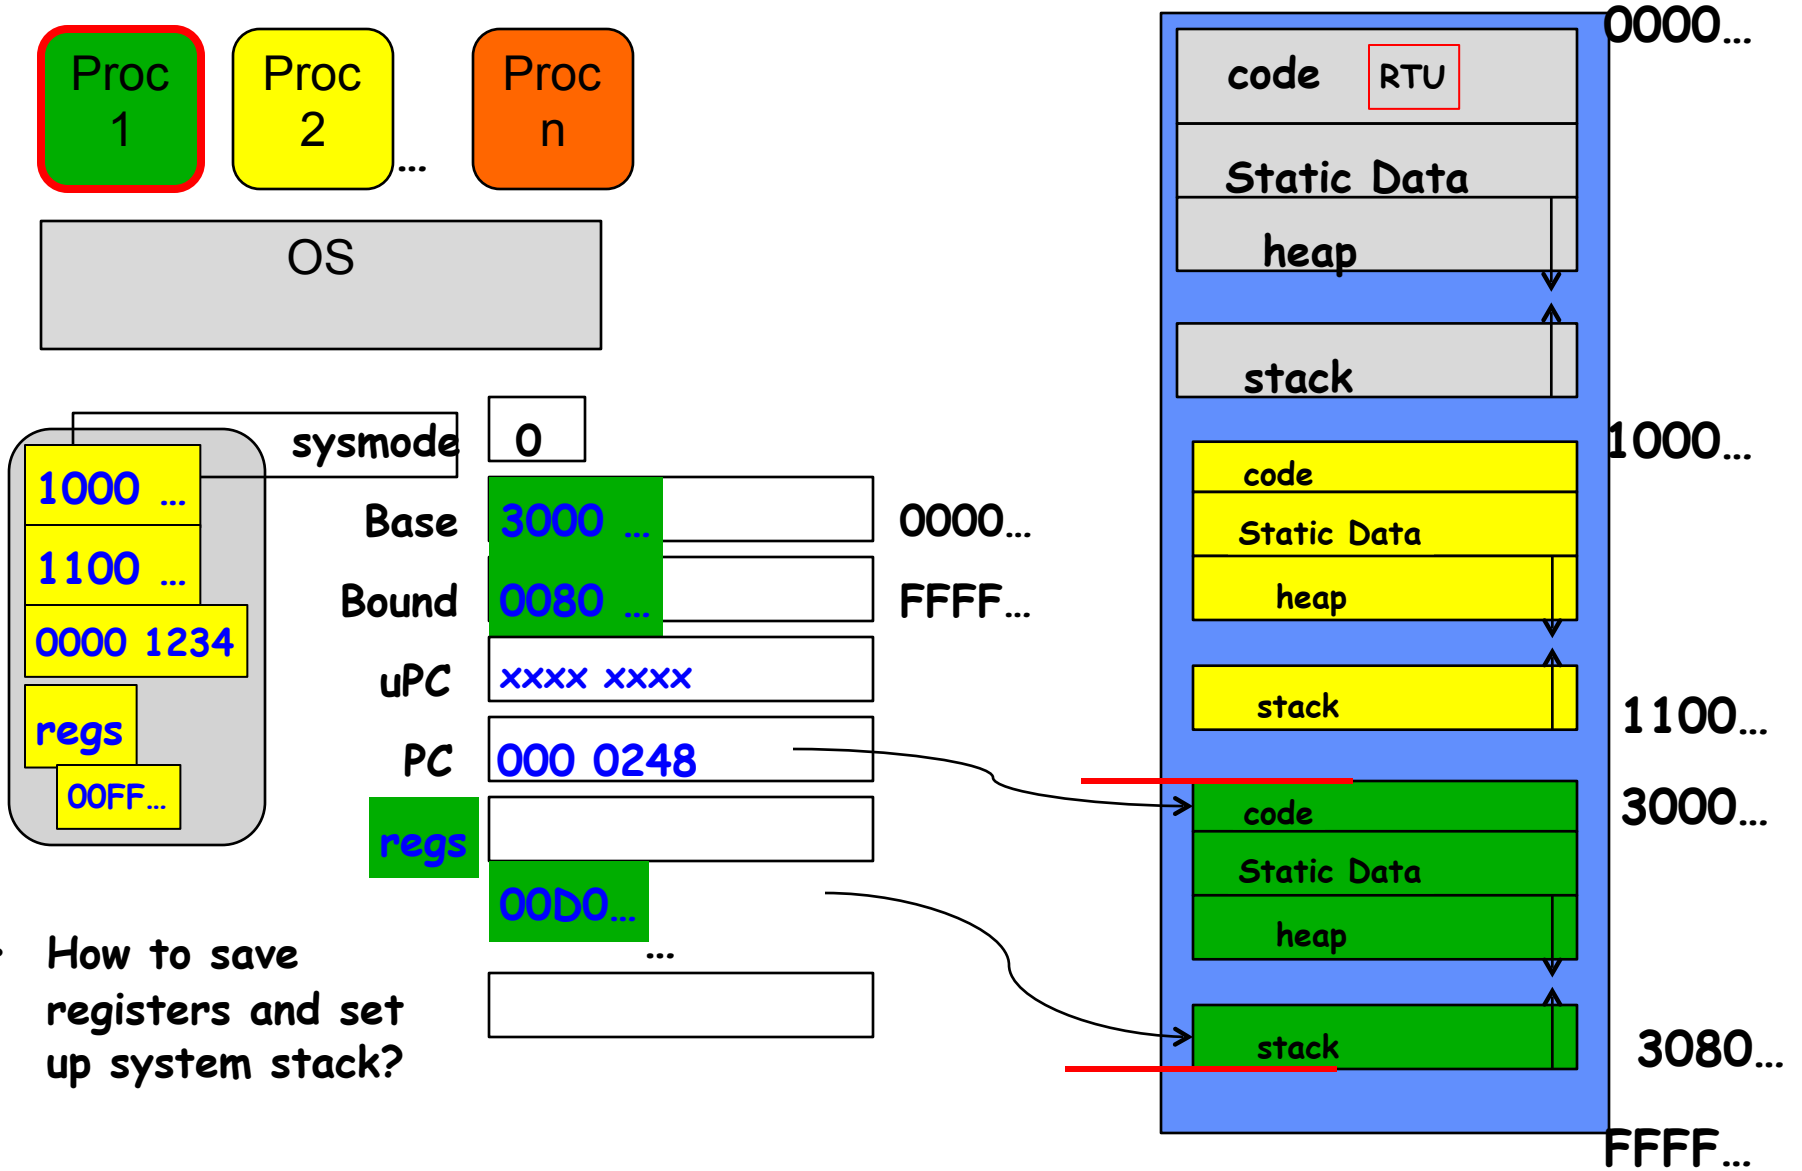
# Simple B&B: Switch User Process



- How to save registers and set up system stack?



# Simple B&B: "resume"




- How to save registers and set up system stack?

---

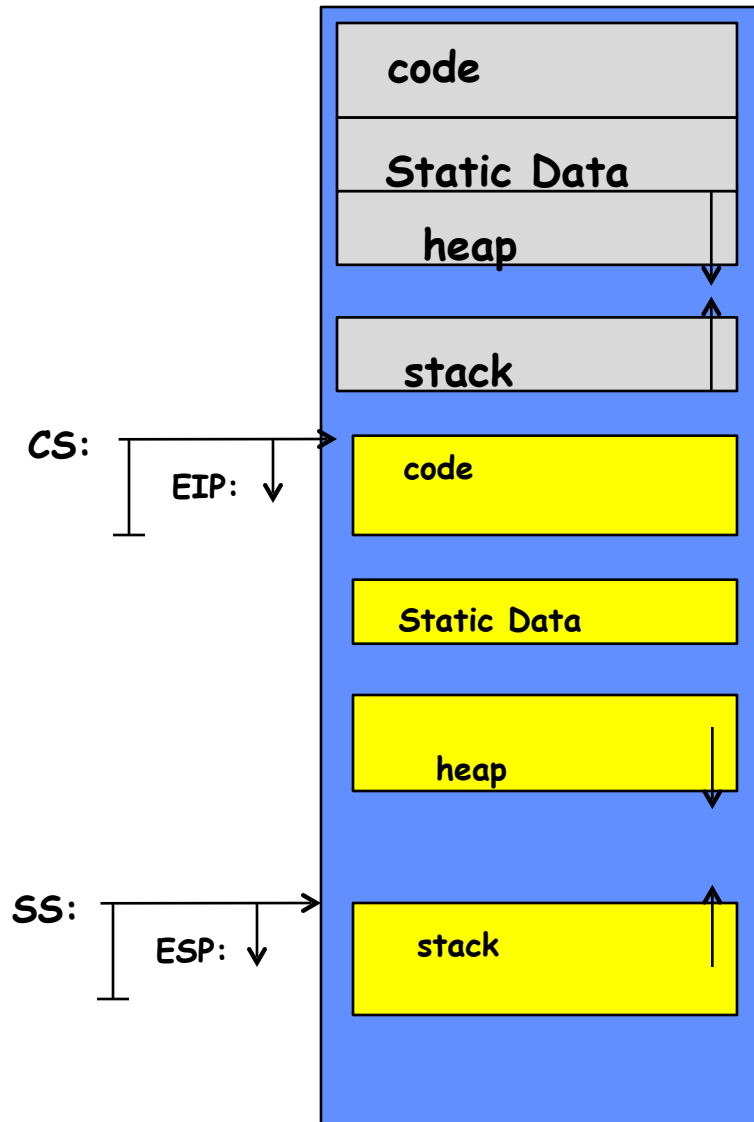
What's wrong with this simplistic address translation mechanism?

# x86 - segments and stacks

## Processor Registers



Start address, length  
and access rights  
associated with each  
segment

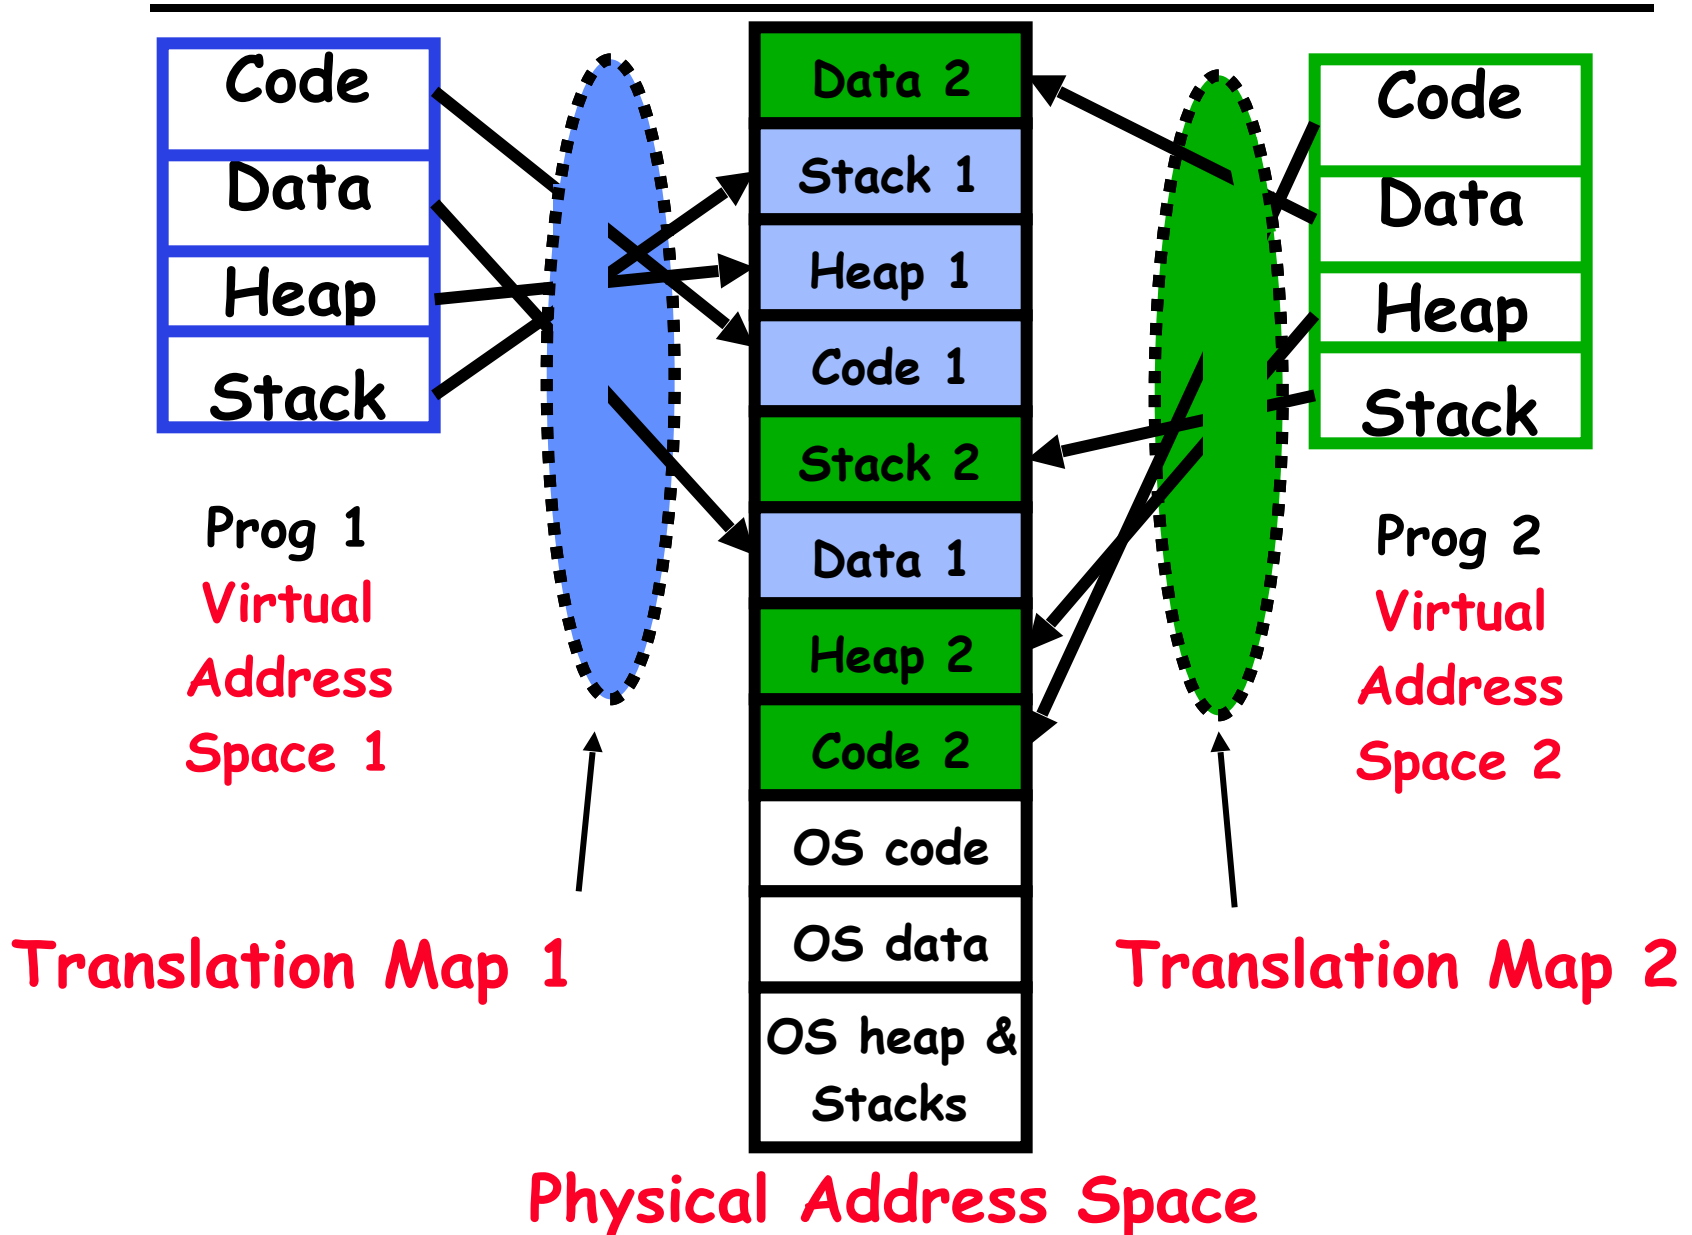


# Virtual Address Translation

---

- **Simpler, more useful schemes too!**
- **Give every process the illusion of its own BIG FLAT ADDRESS SPACE**
  - Break it into pages
  - More on this later

Providing Illusion of Separate Address Space:  
Load new Translation Map on Switch



# Conclusion: Four fundamental OS concepts

---

- **Thread**
  - Single unique execution context
  - Program Counter, Registers, Execution Flags, Stack
- **Address Space w/ Translation**
  - Programs execute in an address space that is distinct from the memory space of the physical machine
- **Process**
  - An instance of an executing program is a process consisting of an address space and one or more threads of control
- **Dual Mode operation/Protection**
  - Only the "system" has the ability to access certain resources
  - The OS and the hardware are protected from user programs and user programs are isolated from one another by controlling the translation from program virtual addresses to machine physical addresses