


MINIMUM	MEDIAN	MAXIMUM	MEAN	STD DEV
30.81	63.23	97.07	63.28	11.62

Prelim 1: Final results

Living dangerously: Safe, Unsafe, Deadlocked



A system's trajectory
through its state space

- **Safe:** For any possible set of resource requests, there exists one **safe schedule** of processing requests that succeeds in granting all pending and future requests
 - no deadlock as long as system can **enforce** that safe schedule!
- **Unsafe:** There exists a set of (pending and future) resource requests that leads to a deadlock, independent of the schedule in which requests are processed
 - unlucky set of requests can force deadlock
- **Deadlocked:** The system has at least one deadlock

The Banker's books

- Assume n processes, m resources
- Max_{ij} = max amount of units of resource R_j needed by P_i
 - MaxClaim_i : Vector of size m such that $\text{MaxClaim}_i[j] = \text{Max}_{ij}$
- Holds_{ij} = current allocation of R_j held by P_i
 - HasNow_i = Vector of size m such that $\text{HasNow}_i[j] = \text{Holds}_{ij}$
- Available = Vector of size m such that $\text{Available}[j]$ = units of R_j available
- A request by P_k is safe if, assuming the request is granted, there is a permutation of P_1, P_2, \dots, P_n such that, for all P_i in the permutation

$$\text{Needs}_i = \text{MaxClaim}_i - \text{HasNow}_i \leq \text{Avail} + \sum_{j=1}^{i-1} \text{HasNow}_j$$

An Example

- 5 processes, 4 resources

Max					Holds					Available				
	R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄
P ₁	0	0	1	2	P ₁	0	0	1	2		1	5	2	0
P ₂	1	7	5	0	P ₂	1	0	0	0					
P ₃	2	3	5	6	P ₃	1	3	5	3					
P ₄	0	6	5	2	P ₄	0	6	3	2					
P ₅	0	6	5	6	P ₅	0	0	1	4					

- Is this a safe state?

An Example

- 5 processes, 4 resources

Max					Holds					Available					Needs				
P ₁	R ₂	R ₃	R ₄		P ₁	R ₁	R ₂	R ₃	R ₄	R ₁	R ₂	R ₃	R ₄		P ₁	R ₁	R ₂	R ₃	R ₄
P ₁	0	0	1	2	P ₁	0	0	1	2	1	5	2	0		P ₁	0	0	0	0
P ₂	1	7	5	0	P ₂	1	0	0	0						P ₂	0	7	5	0
P ₃	2	3	5	6	P ₃	1	3	5	3						P ₃	1	0	0	3
P ₄	0	6	5	2	P ₄	0	6	3	2						P ₄	0	0	2	0
P ₅	0	6	5	6	P ₅	0	0	1	4						P ₅	0	6	4	2

- Is this a safe state?

P₁, P₄, P₂, P₃, P₅

- While safe permutation does not include all processes:
 - Is there a P_i such that Needs_i ≤ Avail?
 - if no, exit with **unsafe**
 - if yes, add P_i to the sequence and set Avail = Avail + HasNow_i
- Exit with **safe**

An Example

- 5 processes, 4 resources

Max					Holds					Available					Needs				
	R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄
P ₁	0	0	1	2	P ₁	0	0	1	2	1	5	2	0	P ₁	0	0	0	0	
P ₂	1	7	5	0	P ₂	1	0	0	0		P ₂	0	7	5	0				
P ₃	2	3	5	6	P ₃	1	3	5	3		P ₃	1	0	0	3				
P ₄	0	6	5	2	P ₄	0	6	3	2		P ₄	0	0	2	0				
P ₅	0	6	5	6	P ₅	0	0	1	4		P ₅	0	6	4	2				

- P₂ wants to change its holdings to 0 4 2 0

An Example

- 5 processes, 4 resources

Max					Holds					Available					Needs				
	R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄
P ₁	0	0	1	2	P ₁	0	0	1	2		2	1	0	0	P ₁	0	0	0	0
P ₂	1	7	5	0	P ₂	0	4	2	0						P ₂	1	3	3	0
P ₃	2	3	5	6	P ₃	1	3	5	3						P ₃	1	0	0	3
P ₄	0	6	5	2	P ₄	0	6	3	2						P ₄	0	0	2	0
P ₅	0	6	5	6	P ₅	0	0	1	4						P ₅	0	6	4	2

- P₂ wants to change its holdings to 0 4 2 0
- Safe? Reduce P₁

An Example

- 5 processes, 4 resources

Max					Holds					Available					Needs				
	R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄		R ₁	R ₂	R ₃	R ₄
P ₁	0	0	0	0	P ₁	0	0	0	0		2	1	1	2	P ₁	0	0	0	0
P ₂	1	7	5	0	P ₂	0	4	2	0						P ₂	1	3	3	0
P ₃	2	3	5	6	P ₃	1	3	5	3						P ₃	1	0	0	3
P ₄	0	6	5	2	P ₄	0	6	3	2						P ₄	0	0	2	0
P ₅	0	6	5	6	P ₅	0	0	1	4						P ₅	0	6	4	2

- P₂ wants to change its holdings to 0 4 2 0
- Safe? Reduce P₁; can't reduce any further

Unsafe!

If all processes were to ask together all the resources they may need, **deadlock!**

Reactive Responses to Deadlock

• Deadlock Detection

- Track resource allocation (who has what)
- Track pending requests (who's waiting for what)

• When should it run?

- For each request?
- After each unsatisfiable request?
- Every hour?
- Once CPU utilization drops below a threshold?

Detecting Deadlock

- 5 processes, 3 resources.

	Holds		
	R ₁	R ₂	R ₃
P ₁	0	1	0
P ₂	2	0	0
P ₃	3	0	3
P ₄	2	1	1
P ₅	0	0	2

Available		
R ₁	R ₂	R ₃
0	0	0

	Pending		
	R ₁	R ₂	R ₃
P ₁	0	0	0
P ₂	2	0	2
P ₃	0	0	0
P ₄	1	0	2
P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	1	0		0	0	0	P ₁	0	0	0
P ₂	2	0	0					P ₂	2	0	2
→ P ₃	3	0	3					P ₃	0	0	0
P ₄	2	1	1					P ₄	1	0	2
P ₅	0	0	2					P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	1	0		3	0	3	P ₁	0	0	0
P ₂	2	0	0					P ₂	2	0	2
→ P ₃	0	0	0					P ₃	0	0	0
P ₄	2	1	1					P ₄	1	0	2
P ₅	0	0	2					P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

		Holds			Available			Pending			
		R ₁	R ₂	R ₃	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
→	P ₁	0	1	0	3	0	3	P ₁	0	0	0
	P ₂	2	0	0				P ₂	2	0	2
	P ₃	0	0	0				P ₃	0	0	0
	P ₄	2	1	1				P ₄	1	0	2
	P ₅	0	0	2				P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

		Holds			Available			Pending			
		R ₁	R ₂	R ₃	R ₁	R ₂	R ₃	R ₁	R ₂	R ₃	
→	P ₁	0	0	0	3	1	3	P ₁	0	0	0
	P ₂	2	0	0				P ₂	2	0	2
	P ₃	0	0	0				P ₃	0	0	0
	P ₄	2	1	1				P ₄	1	0	2
	P ₅	0	0	2				P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	0	0		3	1	3	P ₁	0	0	0
P ₂	2	0	0					P ₂	2	0	2
P ₃	0	0	0					P ₃	0	0	0
→ P ₄	2	1	1					P ₄	1	0	2
P ₅	0	0	2					P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	0	0		5	2	4	P ₁	0	0	0
P ₂	2	0	0					P ₂	2	0	2
P ₃	0	0	0					P ₃	0	0	0
→ P ₄	0	0	0					P ₄	0	0	0
P ₅	0	0	2					P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	0	0		5	2	4	P ₁	0	0	0
P ₂	2	0	0					P ₂	2	0	2
P ₃	0	0	0					P ₃	0	0	0
P ₄	0	0	0					P ₄	0	0	0
P ₅	0	0	2					P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	0	0		7	2	4	P ₁	0	0	0
P ₂	0	0	0					P ₂	0	0	0
P ₃	0	0	0					P ₃	0	0	0
P ₄	0	0	0					P ₄	0	0	0
P ₅	0	0	2					P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	0	0		7	2	4	P ₁	0	0	0
P ₂	0	0	0					P ₂	0	0	0
P ₃	0	0	0					P ₃	0	0	0
P ₄	0	0	0					P ₄	0	0	0
→ P ₅	0	0	2					P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	0	0		7	2	6	P ₁	0	0	0
P ₂	0	0	0					P ₂	0	0	0
P ₃	0	0	0					P ₃	0	0	0
P ₄	0	0	0					P ₄	0	0	0
→ P ₅	0	0	0					P ₅	0	0	0

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Yes, there
is a safe
schedule!

Detecting Deadlock

- 5 processes, 3 resources.

	Holds		
	R ₁	R ₂	R ₃
P ₁	0	1	0
P ₂	2	0	0
P ₃	3	0	3
P ₄	2	1	1
P ₅	0	0	2

Available		
R ₁	R ₂	R ₃
0	0	0

	Pending		
	R ₁	R ₂	R ₃
P ₁	0	0	0
P ₂	2	0	2
P ₃	0	0	0
P ₄	1	0	2
P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Yes, there
is a safe
schedule!

but it is not a safe state!

Detecting Deadlock

- 5 processes, 3 resources.

	Holds		
	R ₁	R ₂	R ₃
P ₁	0	1	0
P ₂	2	0	0
P ₃	3	0	3
P ₄	2	1	1
P ₅	0	0	2

Available		
R ₁	R ₂	R ₃
0	0	0

	Pending		
	R ₁	R ₂	R ₃
P ₁	0	0	0
P ₂	2	0	2
P ₃	0	0	1
P ₄	1	0	2
P ₅	0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- But can determine if the state has a deadlock
 - Given the set of pending requests, is there a safe sequence?
If no, deadlock

Detecting Deadlock

- 5 processes, 3 resources.

	Holds				Available				Pending		
	R ₁	R ₂	R ₃		R ₁	R ₂	R ₃		R ₁	R ₂	R ₃
P ₁	0	1	0		0	0	0		0	0	0
P ₂	2	0	0						2	0	2
P ₃	3	0	3						0	0	1
P ₄	2	1	1						1	0	2
P ₅	0	0	2						0	0	2

- Cannot determine whether the state is safe
 - I need **Max** and **Needs** for that!
- Without Max, can we avoid deadlock by delaying granting requests?
 - NO!** Deadlock triggered when request formulated, not granted!

Deadlock Recovery

- Blue screen & reboot
- Kill one/all deadlocked processes
 - Pick a victim (how?); Terminate; Repeat as needed
 - ▶ Can leave system in inconsistent state
- Proceed without the resource (if application permits)
 - Example: timeout on inventory check at Amazon
- Use transactions
 - Rollback & Restart
 - Need to pick a victim...

Summary

👁 Prevent

- ❑ Negate one of the four necessary conditions

👁 Avoid

- ❑ Schedule processes carefully

👁 Detect

- ❑ Has a deadlock occurred?

👁 Recover

- ❑ Kill or Rollback


Barrier Synchronization

And now for something completely different

- Remember **Mutual Exclusion**...
 - at most one process executing at the same time a given piece of code
- **Barrier Synchronization** is almost the opposite...
 - A set of processes that runs in rounds
 - All must complete the current round before any can start the next
 - Popular in HPC, simulations, graph processing, model checking...

Barrier Abstraction

- **Barrier(\mathcal{N}):** barrier for \mathcal{N} threads
- **bwait(?barr):** wait for every thread to catch up on barrier barr



Test Program for Barriers

```
1  import barrier
2
3  const NTHREADS = 3
4  const NROUNDS = 4
5
6  round = [0,] * NTHREADS
7  invariant (max(round) - min(round)) <= 1
8
9  barr = barrier.Barrier(NTHREADS)
10
11 def thread(self):
12     for r in {0..NROUNDS-1}:
13         barrier.bwait(?barr)
14         round[self] += 1
15
16 for i in {0..NTHREADS-1}:
17     spawn thread(i)
```

*Threads can
be at most
one round
apart*

*Waiting for
everyone to
reach the
barrier*

Barrier Implementation

```
1  from synch import *
2
3  def Barrier(required):
4      result = {
5          .mutex: Lock(), .cond: Condition(),
6          .required: required, .left: required, .cycle: 0
7      }
8
9  def bwait(b):
10     acquire(?b→mutex)
11     b→left -= 1
12     if b→left == 0:
13         b→cycle = (b→cycle + 1) % 2
14         b→left = b→required
15         notifyAll(?b→cond)
16     else:
17         let cycle = b→cycle:
18             while b→cycle == cycle:
19                 wait(?b→cond, ?b→mutex)
20     release(?b→mutex)
```

lock and condition (as you would expect)

no. of threads

*threads that
have not
reached the
barrier*

*incremented
at each round
to allow for
barrier reuse*

*If repeat use were
not an issue, a lock,
a condition variable,
and a counter
initialized to the
number of threads
would suffice,
but,
to reuse the barrier,
we must reset the
counter!*

Using Barriers

```
1  import barrier
2
3  const NTHREADS = 3
4  const NROUNDS = 4
5
6  round = [0,] * NTHREADS
7  invariant (max(round) - min(round)) <= 1
8
9  phase = 0
10 barr = barrier.Barrier(NTHREADS)
11
12 def thread(self):
13     for r in {0..NROUNDS-1}:
14         if self == 0: # coordinator prepares
15             phase += 1
16             barrier.bwait(?barr) # enter parallel work
17             round[self] += 1
18             assert round[self] == phase
19             barrier.bwait(?barr) # exit parallel work
20
21     for i in {0..NTHREADS-1}:
22         spawn thread(i)
```

*everybody waits
for the
coordinator*

*everybody
synchronizes at
the end of round*

Memory Management

(Ch. 12–17)

Abstraction is our Business

- What I have
 - A single (or a finite number) of CPUs
 - Many programs I would like to run
- What I want : a Thread
 - Each program has full control of one or more CPUs

Abstraction is our Business

• What I have

- A certain amount of physical memory
- Multiple programs I would like to run
 - ▶ together, they may need more than the available physical memory

• What I want : **an Address Space**

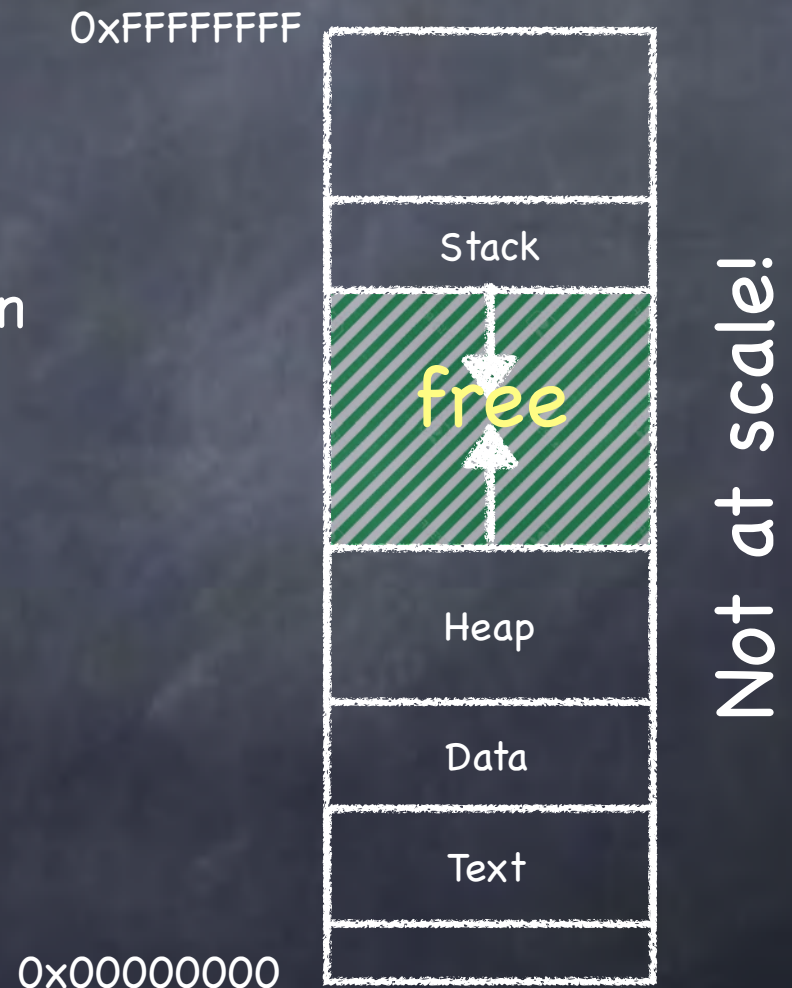
- Each program has as much memory as the machine's architecture will allow to name
- All for itself

Address Space

- Set of all names used to identify and manipulate unique instances of a given resource
 - memory locations (determined by the size of the machine's word)
 - for 32-bit-register machine, the address space goes from 0x00000000 to 0xFFFFFFFF
 - memory locations (determined by the number of memory banks mounted on the machine)
 - phone numbers (XXX) (YYY-YYYY)
 - colors: R (8 bits) + G (8 bits) + B (8 bits)

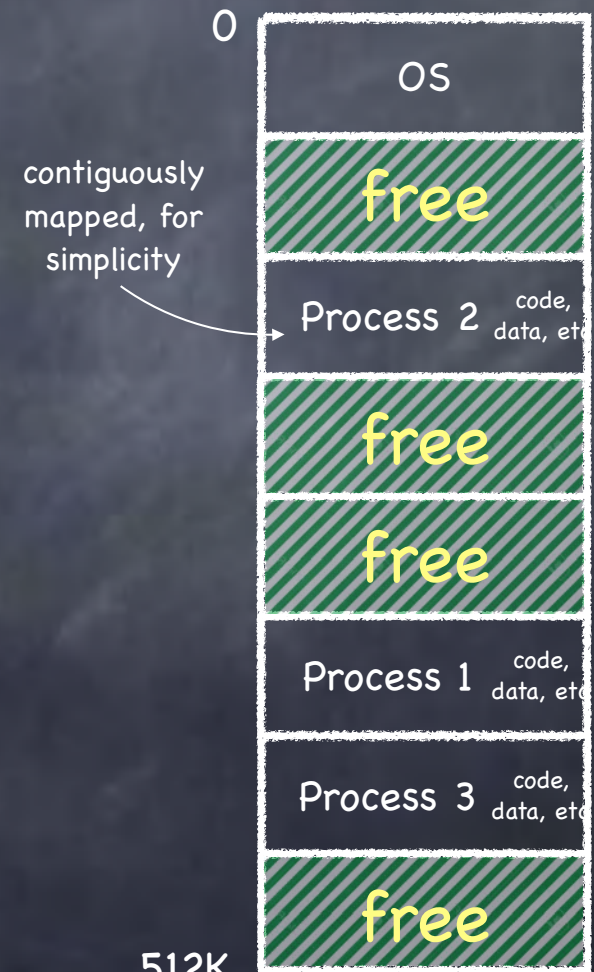
Virtual Address Space: An Abstraction for Memory

- Virtual addresses start at 0
- Heap and stack can be placed far away from each other, so they can nicely grow
- Addresses are all contiguous
- Size is independent of physical memory on the machine



Physical Address Space: How memory actually looks

- Processes loaded in memory at some memory location
 - virtual address 0 is not loaded at physical address 0
- Multiple processes may be loaded in memory at the same time, and yet...
- ...physical memory may be too small to hold even a single virtual address space in its entirety
 - 64-bit, anyone?

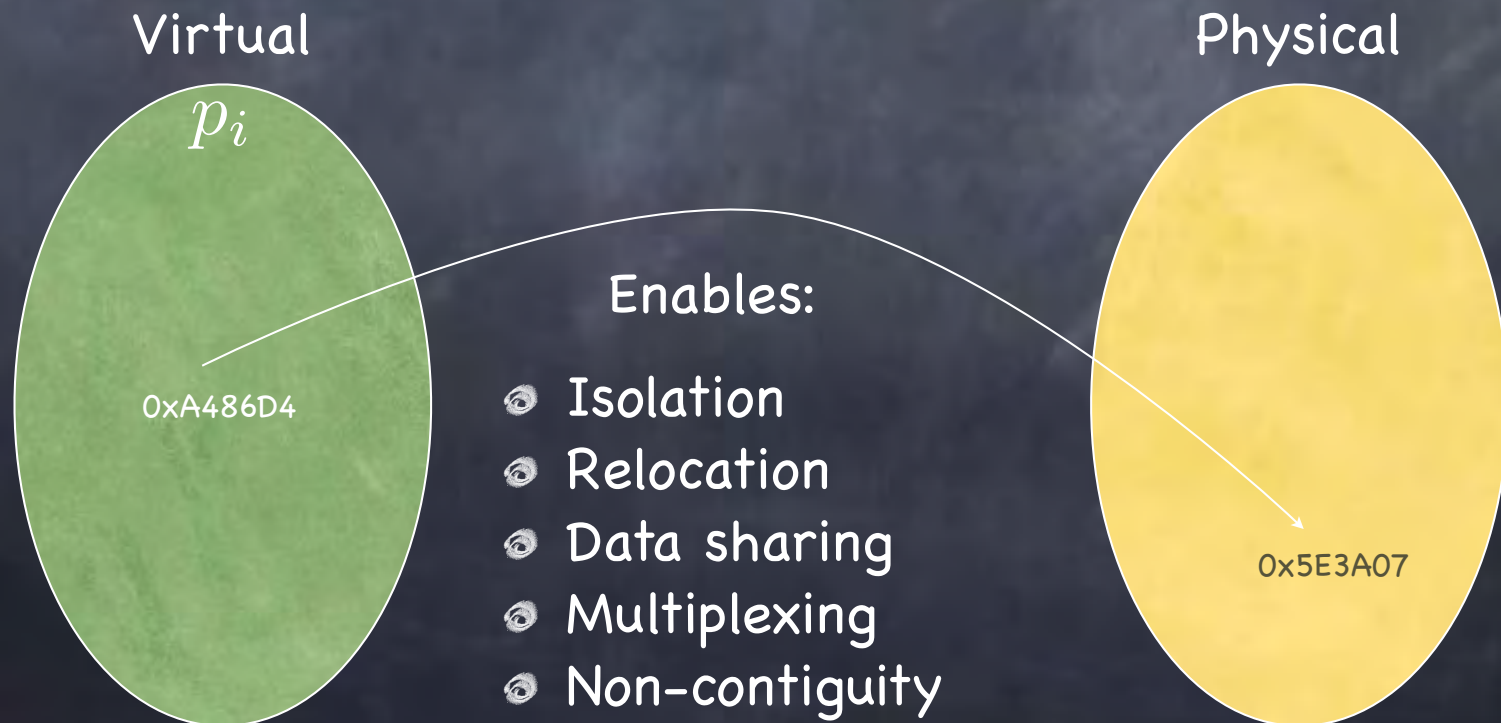


II. Memory Isolation

Step 2: Address Translation

- Implement a function mapping

$\langle pid, virtual\ address \rangle$ into *physical address*



II. Memory Isolation


Step 2: Address Translation

- Implement a function mapping
(pid, virtual address) into physical address



Isolation

- At all times, functions used by different processes map to disjoint ranges — aka “Stay in your room!”



Relocation

- The range of the function used by a process can change over time




Relocation

- The range of the function used by a process can change over time — “Move to a new room!”



Data Sharing

- Map different virtual addresses of distinct processes to the same physical address — (“Share the kitchen”)




Data Sharing

- Map different virtual addresses of distinct processes to the same physical address — (“Share the kitchen”)



Multiplexing

- Create illusion of almost infinite memory by changing domain (set of virtual addresses) that maps to a given range of physical addresses — ever lived in a studio?



Multiplexing

- The domain (set of virtual addresses) that map to a given range of physical addresses can change over time



Multiplexing

- The domain (set of virtual addresses) that map to a given range of physical addresses can change over time



Multiplexing

- The domain (set of virtual addresses) that map to a given range of physical addresses can change over time




Multiplexing

- The domain (set of virtual addresses) that map to a given range of physical addresses can change over time




More Multiplexing

- At different times, different processes can map part of their virtual address space into the same physical memory — (change tenants)



More Multiplexing

- At different times, different processes can map part of their virtual address space into the same physical memory — (change tenants)




(Non) Contiguity

- Contiguous virtual addresses can be mapped to non-contiguous physical addresses...



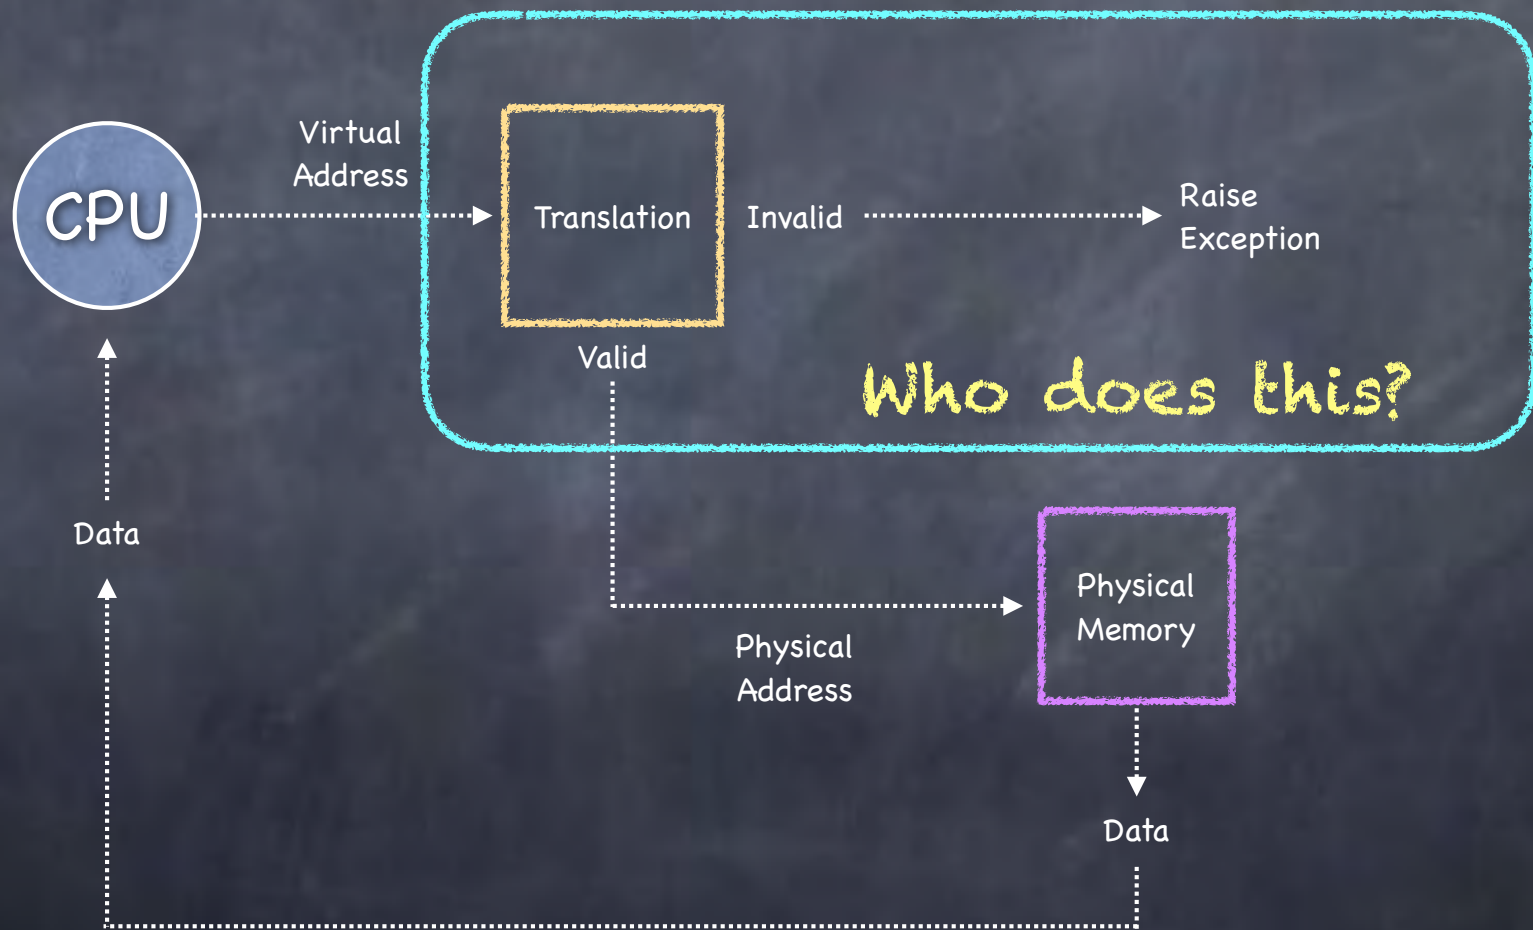
(Non) Contiguity

- ...and non-contiguous virtual addresses can be mapped to contiguous physical addresses



The Power of Mapping

Address Translation, Conceptually



Memory Management Unit (MMU)

- Hardware device
 - Maps virtual addresses to physical addresses
- User process
 - deals with **virtual** addresses
 - never sees the physical address
- Physical memory
 - deals with **physical** addresses
 - never sees the virtual address

