CS111, Lecture 23
Demand Paging

Optional reading:

Operating Systems: Principles and Practice (2nd Edition): Chapter 9

This document is copyright (C) Stanford Computer Science and Nick Troccoli, licensed under Creative Commons Attribution 2.5 License. All rights reserved.
Based on slides and notes created by John Ousterhout, Jerry Cain, Chris Gregg, and others.
NOTICE RE UPLOADING TO WEBSITES: This content is protected and may not be shared, uploaded, or distributed. (without expressed written permission)

Key Questions: how do we make enough space for the page maps? And what happens when we run out of physical pages?

CS111 Topic 4: Virtual Memory

Virtual Memory *- How can one set of memory be shared among several processes? How can the operating system manage access to a limited amount of system memory?*

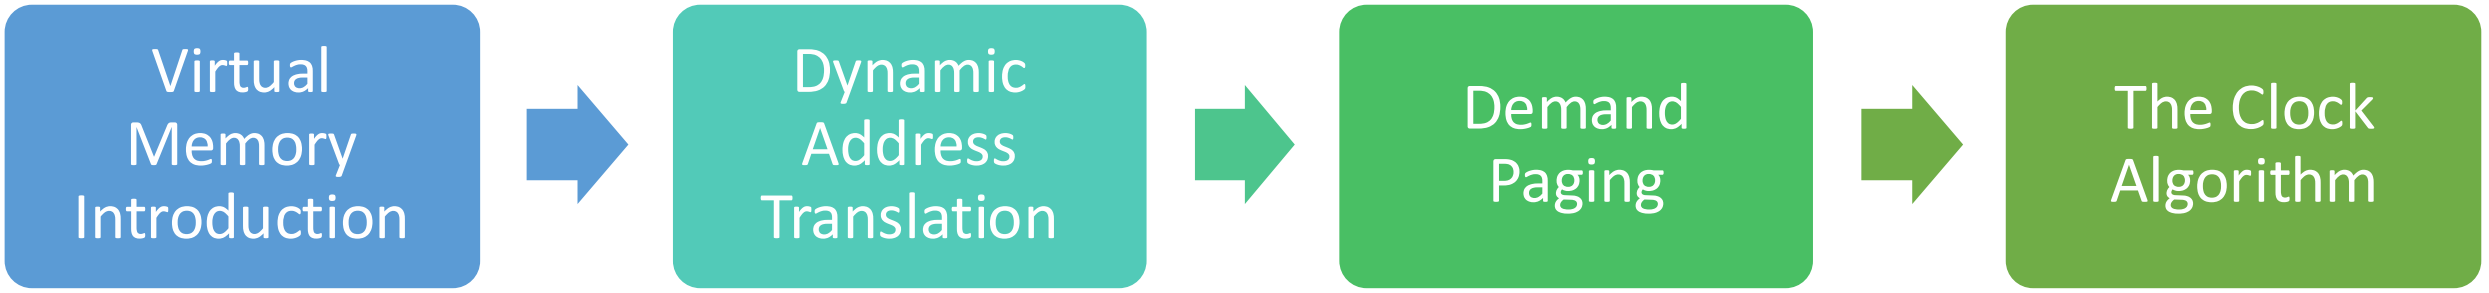

assign6: implement *demand paging* system to translate addresses and load/store memory contents for programs as needed.

Learning Goals

- Learn more about page maps and how they help translate virtual addresses to physical addresses
- Understand how paging allows us to swap memory contents to disk when we need more physical pages.
- Learn about the benefits of demand paging in making memory look larger than it really is

Plan For Today

- **Recap:** Base and bound, multiple segments, and paging
- Page Map Size
- Demand Paging

Plan For Today

- **Recap: Base and bound, multiple segments, and paging**
- Page Map Size
- Demand Paging

Dynamic Address Translation

Key question: how do the MMU / OS translate from virtual addresses to physical ones? Three designs we'll consider:

1. **Base and bound**
2. **Multiple Segments**
3. **Paging**

Approach #2: Multiple Segments

Key Idea: Each process is split among several variable-size areas of memory, called segments.

- E.g. one segment for code, one segment for data/heap, one segment for stack.
- Start of each segment is pinned to a specific virtual address
- OS maps each segment individually – each segment has its own base and bound, kept in a *segment map* for that process (segment map stored in MMU)
- Also store a *protection* bit for each segment: whether the process is allowed to write to it or not in addition to reading
- Now each segment can have its own permissions, grow/shrink independently, be swapped to disk independently, be moved independently, and even be shared between processes (e.g. shared code).
- Top bit(s) of virtual address encode segment number, rest encode offset

Multiple Segments

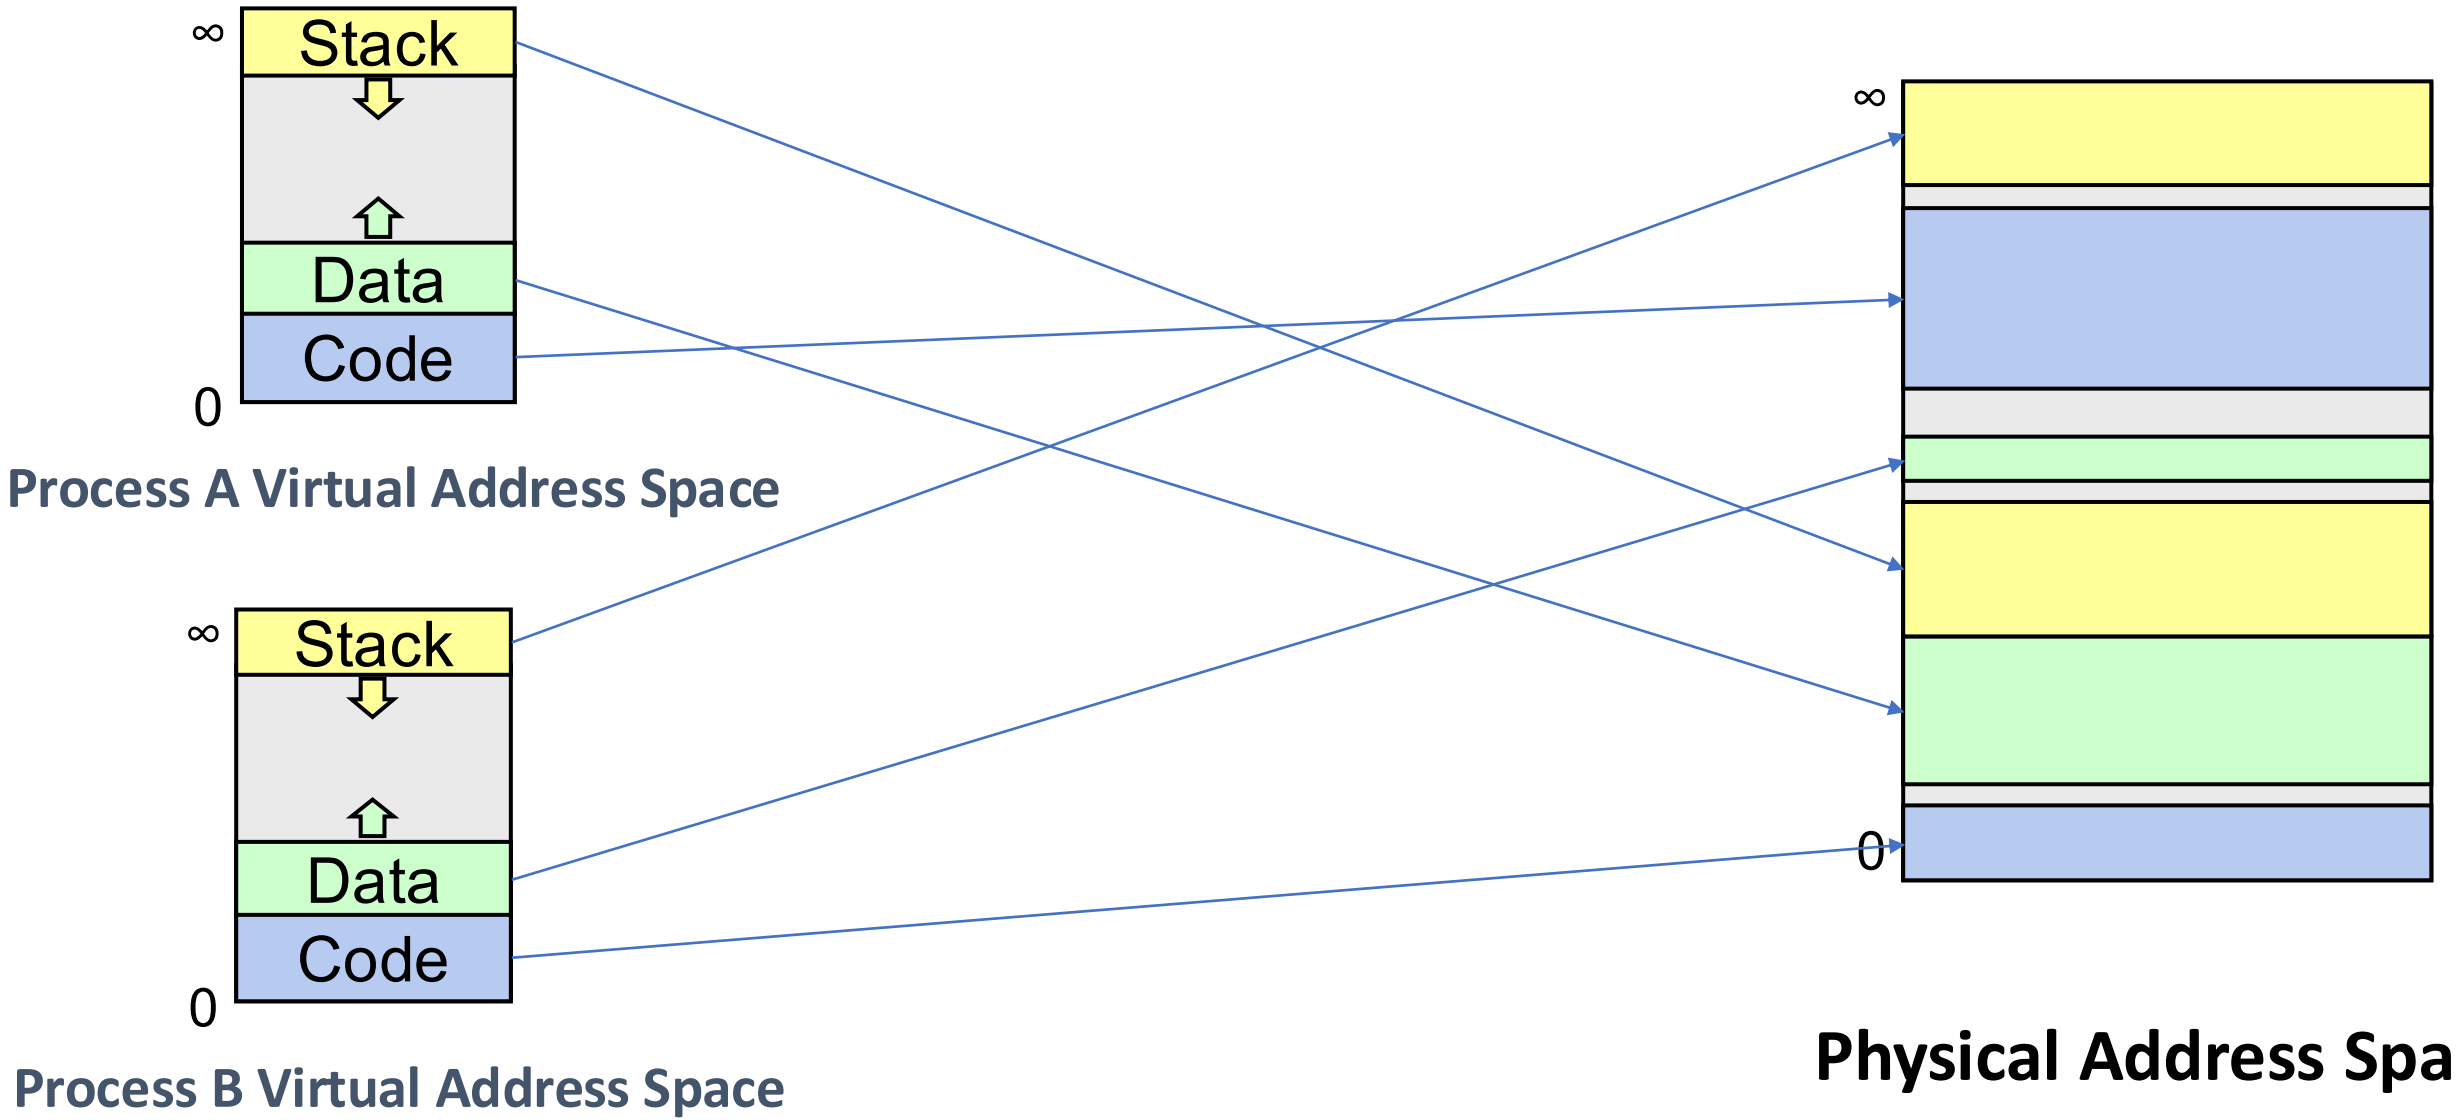

Multiple Segments – Changing A Bound

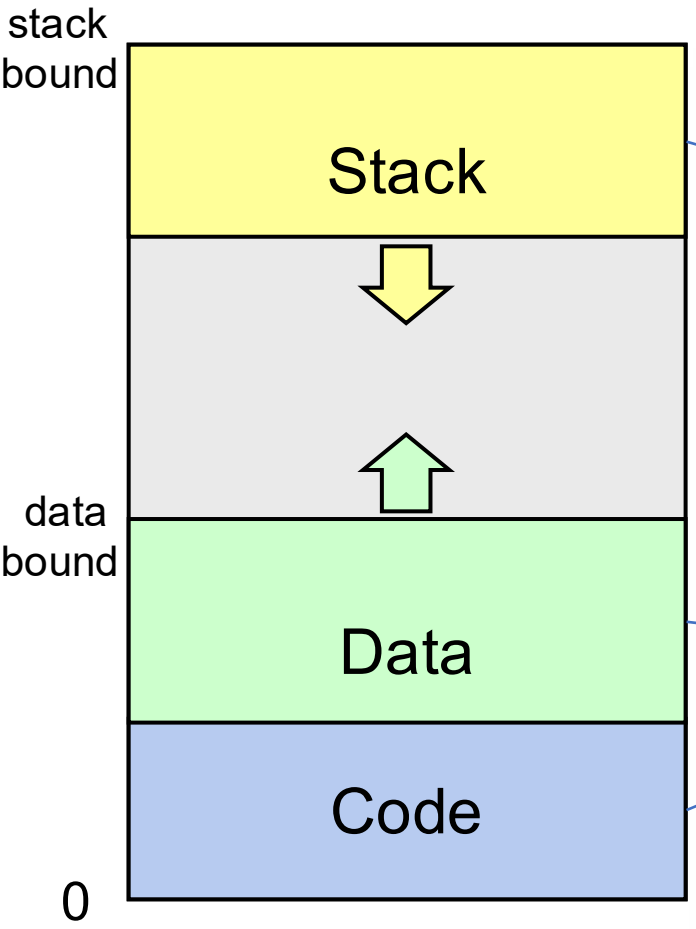

Process A Virtual Address Space

- Buffer space between stack + heap doesn't need to be initially mapped.
- Growing a segment upwards works well for the heap, but not for the stack, for the same reason as base and bound: we can't move existing stack data after the program starts.
- Still fragmentation problem, plus need to decide # bits for segment number vs. offset.

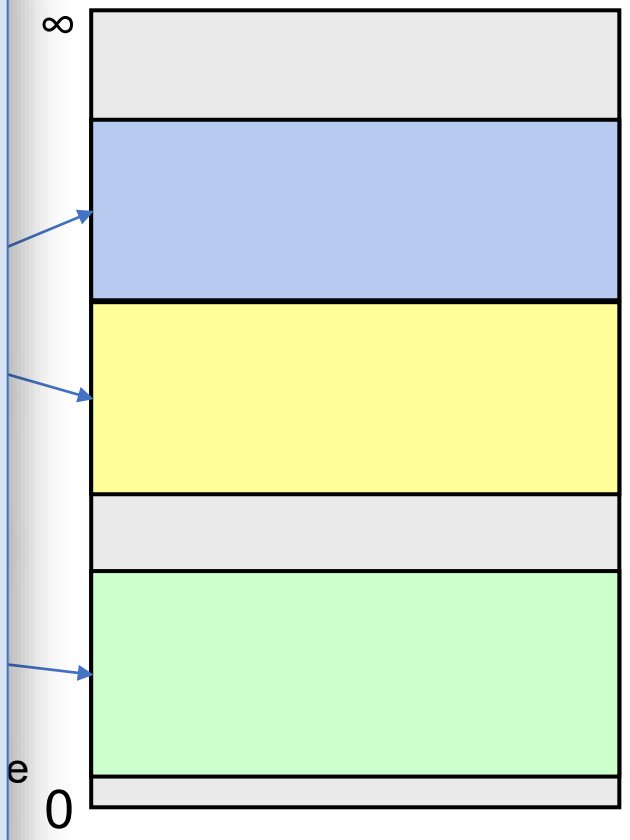

Physical Address Space

Paging

Key Idea: Divide virtual and physical memory into fixed-size chunks called *pages*. (Common size is 4KB pages).

- A "page" of virtual memory maps to a "page" of physical memory. No partial pages. No more external fragmentation! (but some internal fragmentation if not all of a page is used).
- The **page number** is a numerical ID for a page. We have virtual page numbers and physical page numbers.
- Each process has a *page map ("page table")* with an entry for each virtual page, mapping it to a physical page number and other info such as a protection bit (read-only or read-write).
- First bits of an address encode offset, rest encode (virtual or physical) page number

Paging

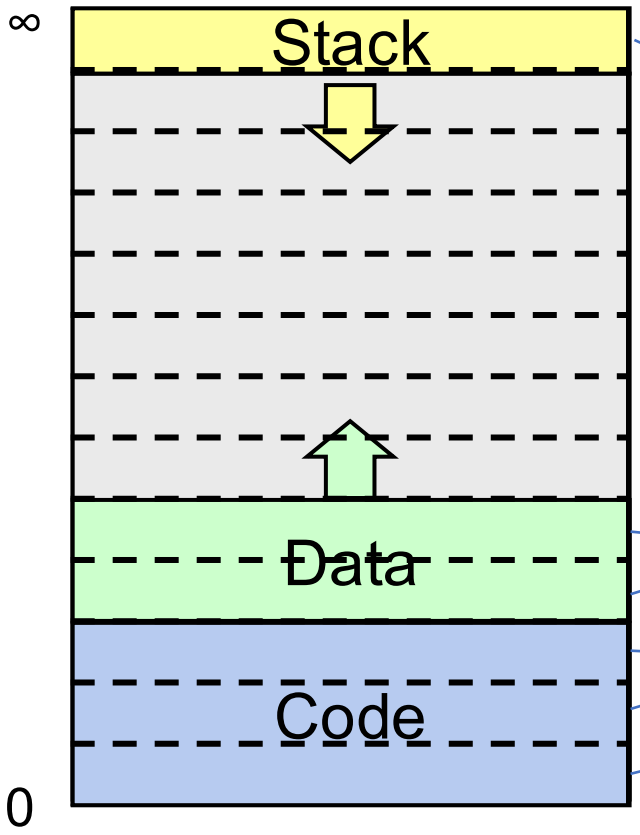

Process A Virtual Address Space

- Do not need to map each segment contiguously. Instead, we map just one page at a time.
- We can later map more pages either up or down, **because the start of the segment is not pinned to a specific address**.
- We can move each page separately in physical memory by modifying the page map.

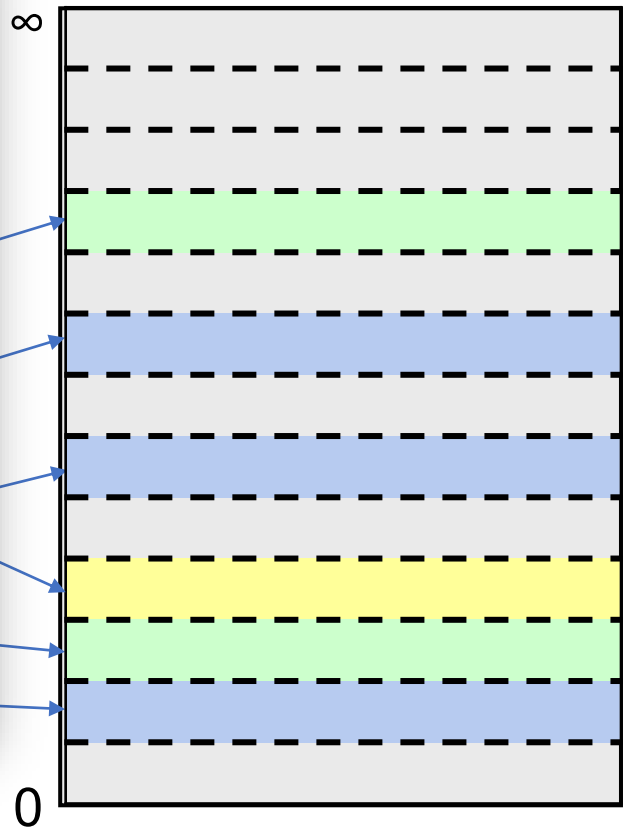

Physical Address Space

Paging

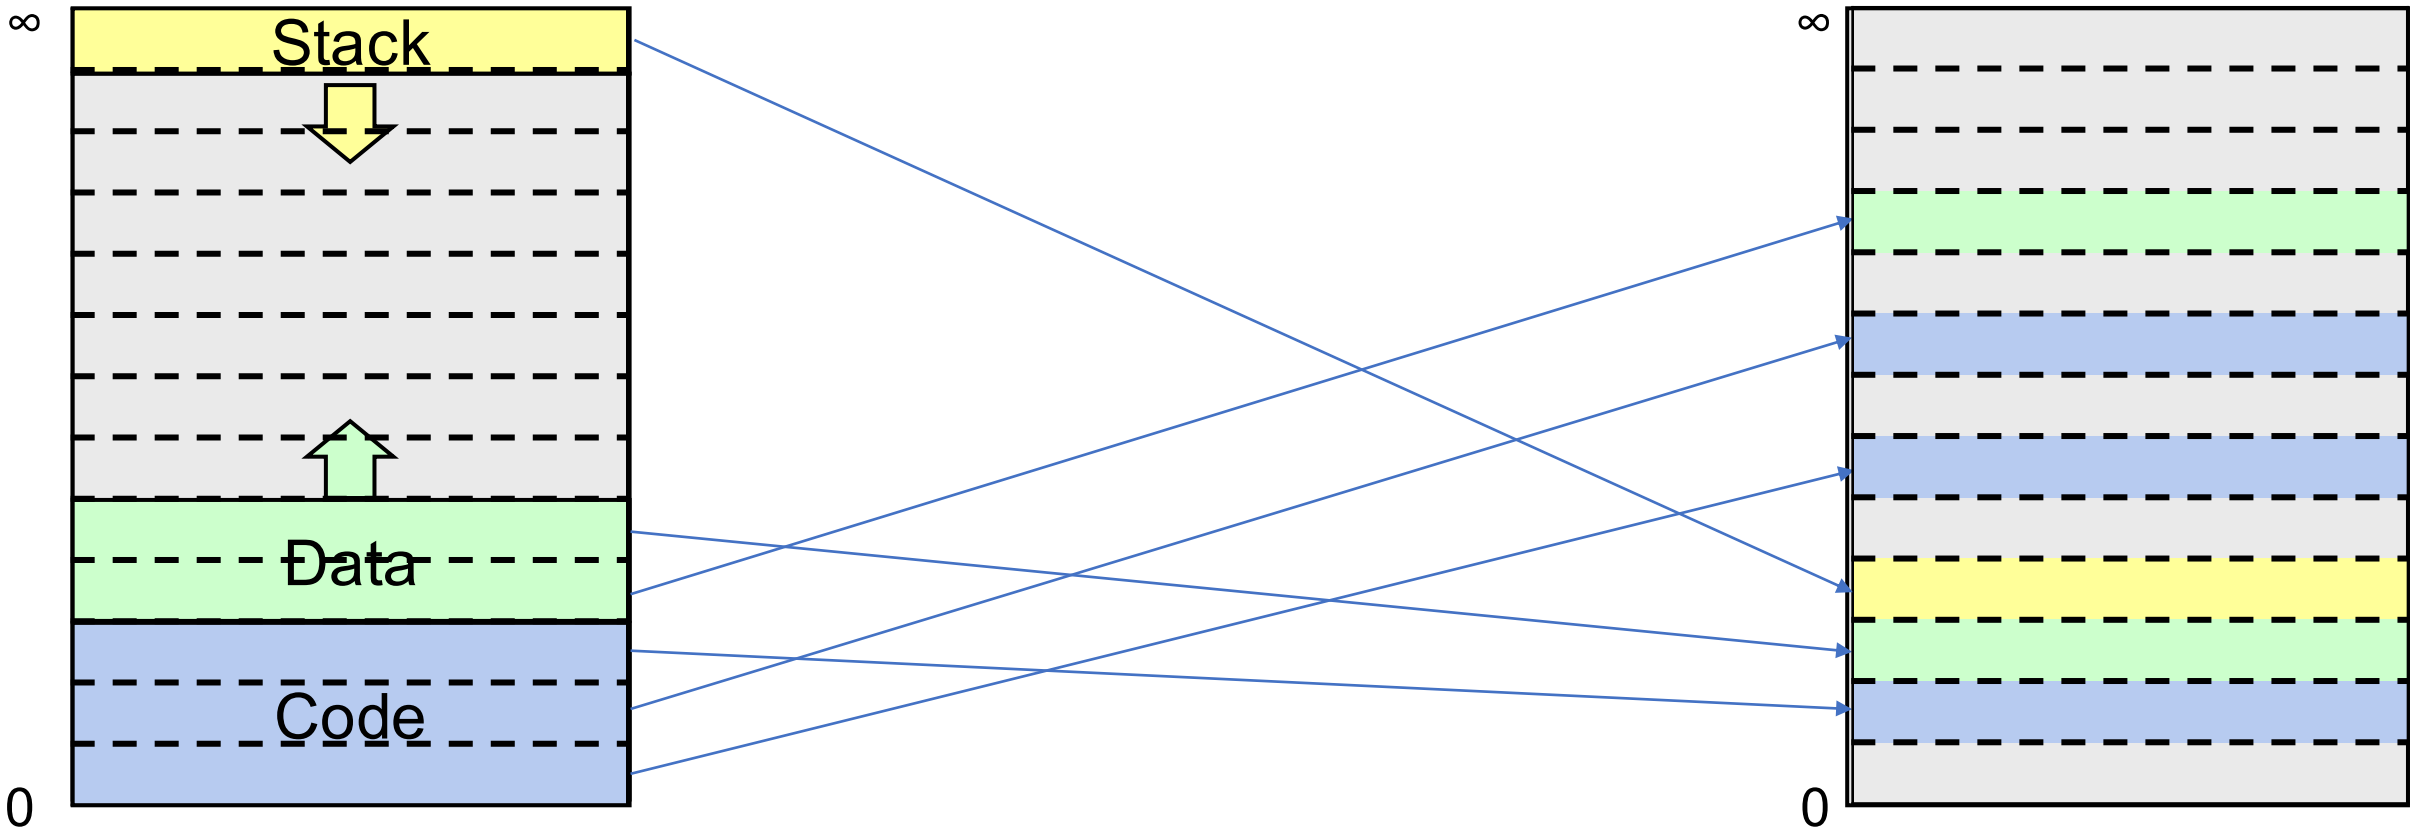

Process A Virtual Address Space

Physical Address Space

Paging

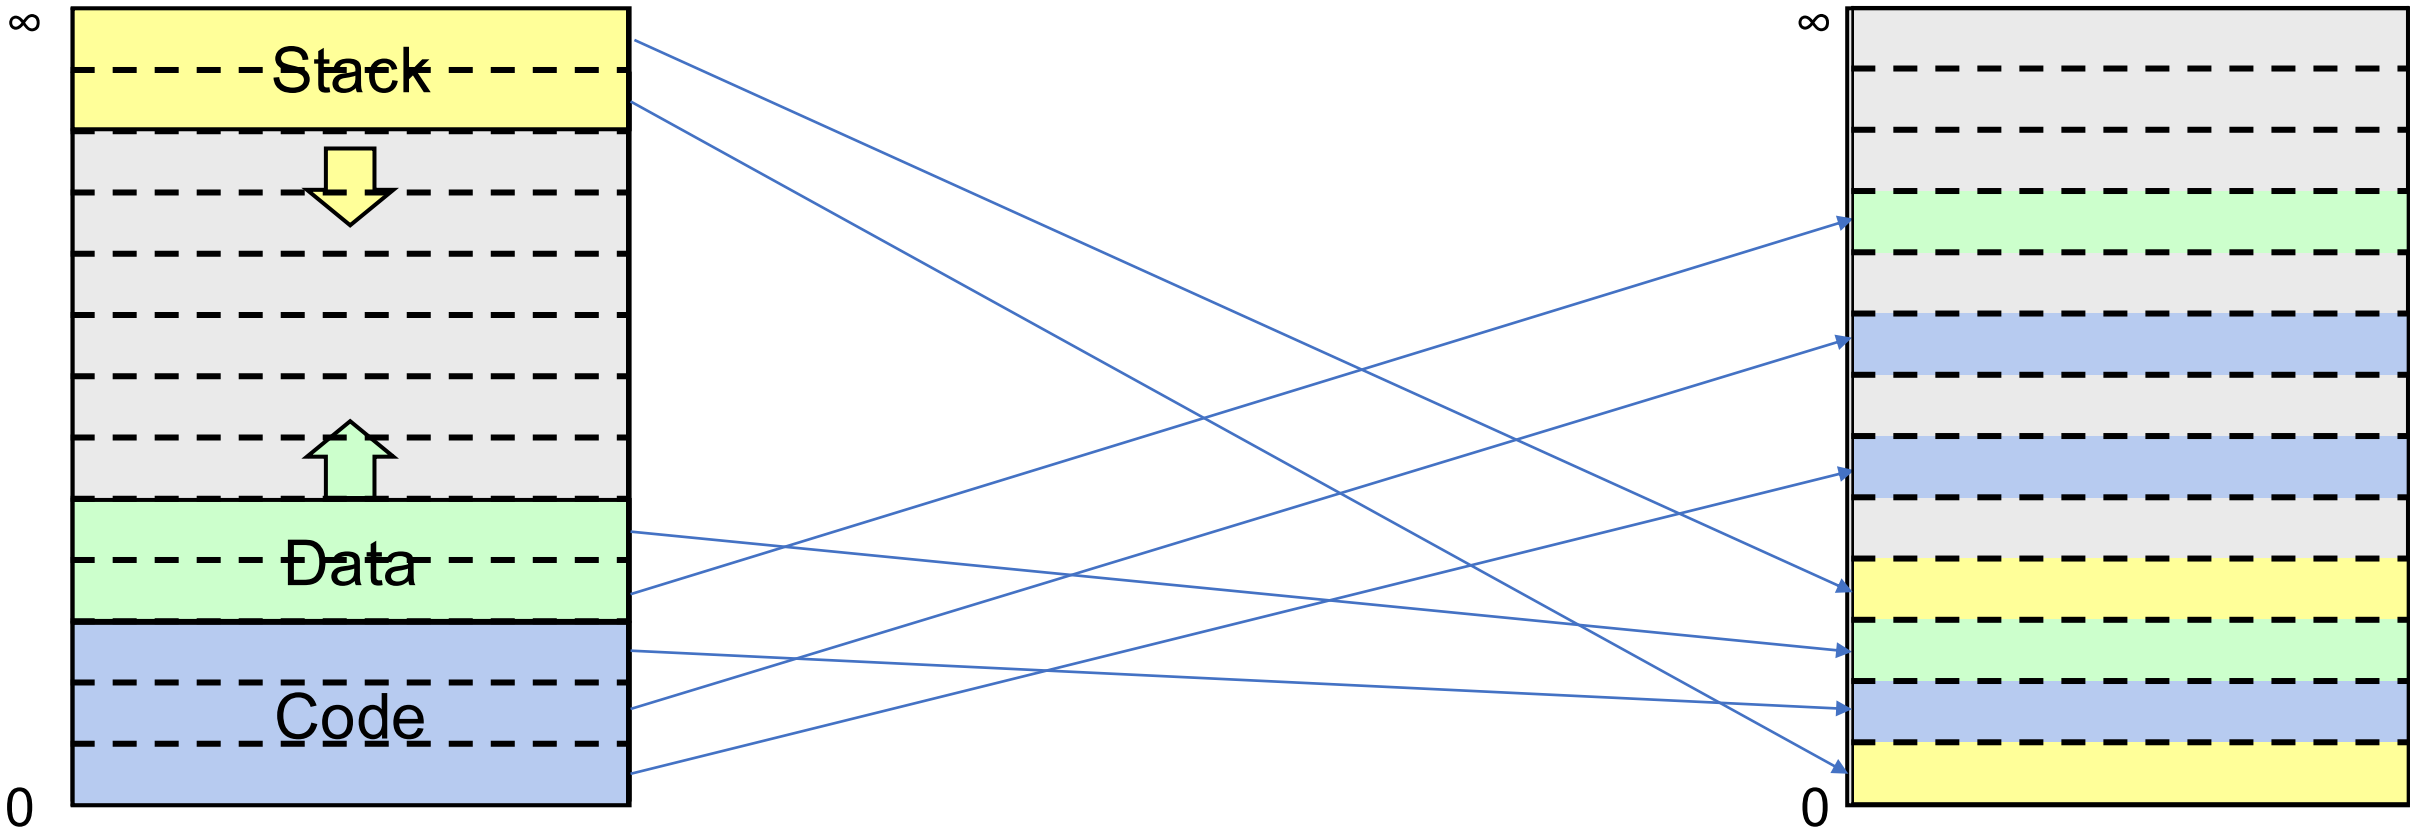

Process A Virtual Address Space

Physical Address Space

Page Map

Index	Physical page #	Writeable?
…	…	…
3	0x2342	1
2	0x12625	1
1	0x13241	0
0	0x256	0



Virtual Address

0x3400

Physical Address

???

Page Map

Index	Physical page #	Writeable?
…	…	…
3	0x2342	1
2	0x12625	1
1	0x13241	0
0	0x256	0

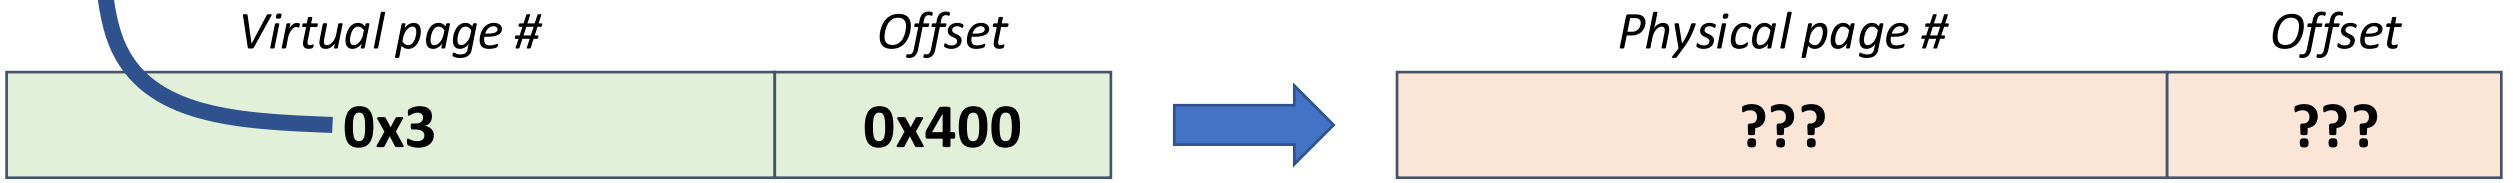

Virtual Address

0x3400

Physical Address

???

Page Map

Index	Physical page #	Writeable?
…	…	…
3	0x2342	1
2	0x12625	1
1	0x13241	0
0	0x256	0

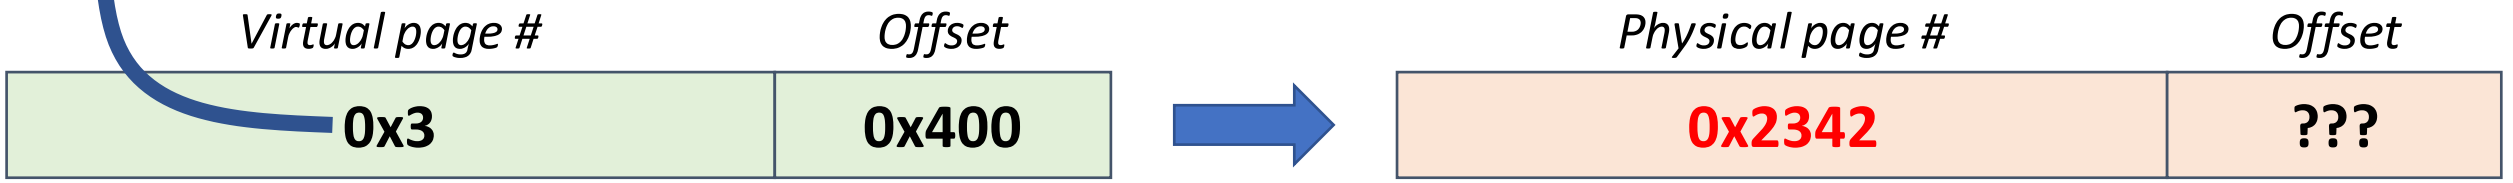

Virtual Address

0x3400

Physical Address

???

Page Map

Index	Physical page #	Writeable?
…	…	…
3	0x2342	1
2	0x12625	1
1	0x13241	0
0	0x256	0

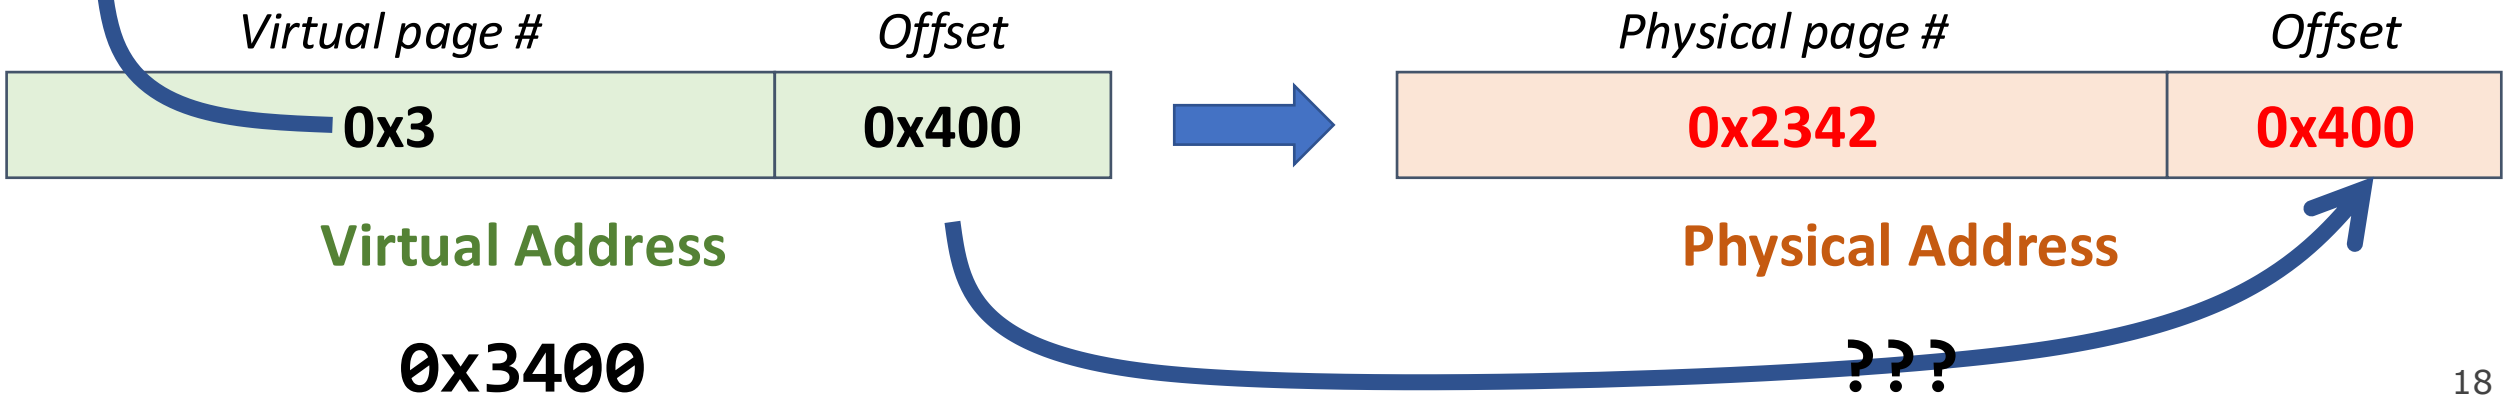

Page Map

Index	Physical page #	Writeable?
…	…	…
3	0x2342	1
2	0x12625	1
1	0x13241	0
0	0x256	0

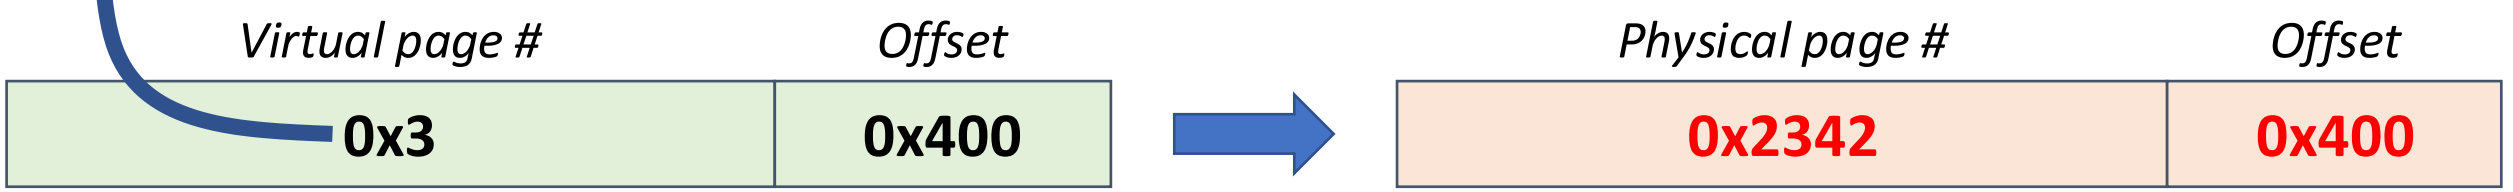

Virtual Address

0x3400

Physical Address

0x2342400

Paging

On each memory reference:

- Look up info for that virtual page in the page map
- If it's a valid virtual page number, get the physical page number it maps to, and combine it with the specified offset to produce the physical address.

Key Idea: to enable quick lookup, the page table is an *array* where the indexes are the virtual page numbers. E.g. index 5 corresponds to virtual page 5.

- (Why not e.g. a hashmap? Not feasible at the time with hardware, plus turns out array ends up working pretty well with optimizations)

<u>This means we must have an entry for every virtual page, even invalid ones</u> – because missing entries would mean indexes wouldn't line up anymore.

Plan For Today

- **Recap:** Base and bound, multiple segments, and paging
- **Page Map Size**
- Demand Paging

Paging

Page map has entries for *all* pages, including invalid ones. We will add an additional field marking whether it's valid ("present").

Page Map

Index	Physical page #	Writeable?	Present?
…	…	…	…
3	0x2342	1	1
2	XXX	X	0
1	0x13241	0	1
0	XXX	X	0

Page Map

Index	Physical page #	Writeable?	Present?
…	…	…	…
3	0x2342	1	1
2	XXX	X	0
1	0x13241	0	1
0	XXX	X	0

If there is a memory access in virtual pages 0 or 2 here, it would trap due to an invalid memory reference.

Page Map Size

Problem: how big is a single process's page map? An entry for *every* page?

Example with x86-64: 36-bit virtual page numbers, 8-byte map entries

How many possible virtual page #s? 2^{36}

2^{36} virtual pages x 8 bytes per page entry = ???

Page Map Size

Problem: how big is a single process's page map? An entry for *every* page?

Example with x86-64: 36-bit virtual page numbers, 8-byte map entries

How many possible virtual page #s? 2^{36}

2^{36} virtual pages x 8 bytes per page entry = **512GB!!** (2^{39} bytes)

Plus, most processes are small, so most pages will be "not present". And even large processes use their address space sparsely (e.g. code at bottom, stack at top).

Page Map Size

x86-64 solution: represent the page map as a multi-level *tree.*

- Top level of page map has entries for *ranges of virtual pages* (0 to 2^{27}-1, 2^{27} to $2 * 2^{27} - 1$, etc.). **Only if** any pages in that range are present does that entry point to a lower level in the tree (saves space). We store the location of the top level to access it.
- Lower levels follow a similar structure – entry for ranges of pages, and they only map to something if at least one of the pages in that range is present.
- The lowest level of the tree contains actual physical page numbers.

Page Map Tree Structure

Tree, level 4:

Virtual page range	Pointer	Present?
…	…	…
$(3 * 2^{27})$ to $(4 * 2^{27} - 1)$	0x2342120	1
$(2 * 2^{27})$ to $(3 * 2^{27} - 1)$	XXXXXX	0
2^{27} to $(2 * 2^{27} - 1)$	XXXXXX	0
0 to $(2^{27} - 1)$	0x423f3210	1

Page Map Tree Structure

Tree, level 4:

Virtual page range	Pointer	Present?
…	…	…
$(3 * 2^{27})$ to $(4 * 2^{27} - 1)$	0x2342120	1
$(2 * 2^{27})$ to $(3 * 2^{27} - 1)$	XXXXXX	0
2^{27} to $(2 * 2^{27} - 1)$	XXXXXX	0
0 to $(2^{27} - 1)$	0x423f3210	1

Page Map Tree Structure

Tree, level 3:

Virtual page range	Pointer	Present?
…	…	…
$(3 * 2^{18})$ to $(4 * 2^{18} - 1)$	0x1692054300	1
$(2 * 2^{18})$ to $(3 * 2^{18} - 1)$	0x45679819f0	1
2^{18} to $(2 * 2^{18} - 1)$	XXXXXX	0
0 to $(2^{18} - 1)$	XXXXXX	0

Page Map Tree Structure

Tree, level 3:

Virtual page range	Pointer	Present?
…	…	…
$(3 * 2^{18})$ to $(4 * 2^{18} - 1)$	0x1692054300	1
$(2 * 2^{18})$ to $(3 * 2^{18} - 1)$	0x45679819f0	1
2^{18} to $(2 * 2^{18} - 1)$	XXXXXX	0
0 to $(2^{18} - 1)$	XXXXXX	0

Page Map Tree Structure

Tree, level 2:

Virtual page range	Pointer	Present?
…	…	…
$(3 * 2^9)$ to $(4 * 2^9 - 1)$	0xff54210	1
$(2 * 2^9)$ to $(3 * 2^9 - 1)$	XXXXXX	0
2^9 to $(2 * 2^9 - 1)$	0xff12900	1
0 to $(2^9 - 1)$	XXXXXX	0

Page Map Tree Structure

Tree, level 2:

Virtual page range	Pointer	Present?
…	…	…
$(3 * 2^9)$ to $(4 * 2^9 - 1)$	0xff54210	1
$(2 * 2^9)$ to $(3 * 2^9 - 1)$	XXXXXX	0
2^9 to $(2 * 2^9 - 1)$	0xff12900	1
0 to $(2^9 - 1)$	XXXXXX	0

Page Map Tree Structure

Tree, level 1:

Index	Physical page	Present?
…	…	…
3	0x4928350	1
2	XXXXXX	0
1	0x9125010	1
0	0x9424020	1

Page Map Tree Structure

Tree, level 1:

Index	Physical page	Present?
…	…	…
3	0x4928350	1
2	XXXXXX	0
1	0x9125010	1
0	0x9424020	1

Page Map Tree Structure

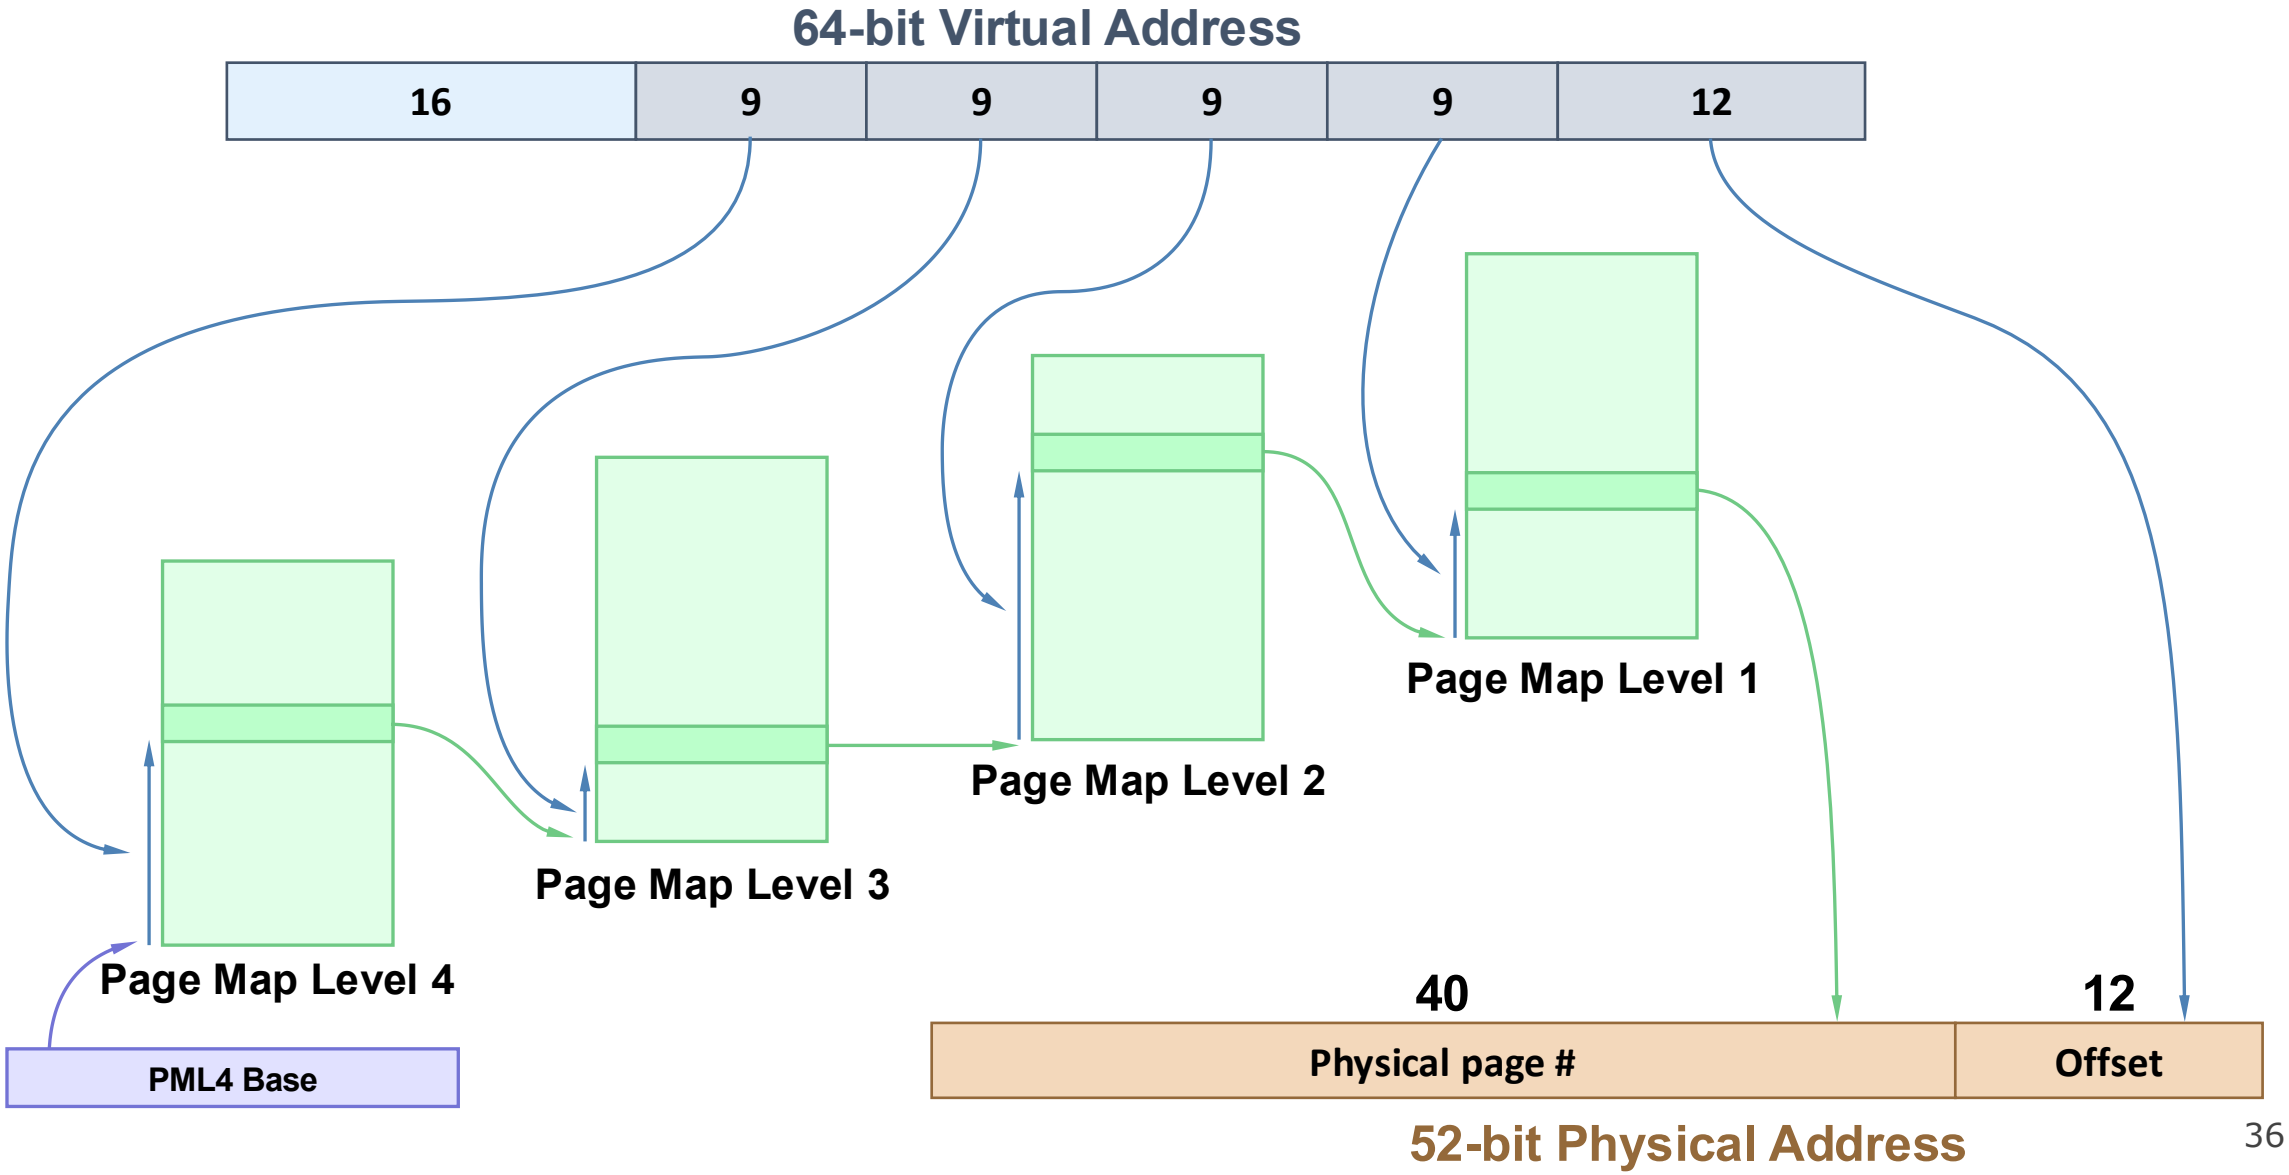

assign6

On assign6, you'll implement your own virtual memory system using paging:

- You'll intercept memory requests
- You'll maintain a page map mapping virtual addresses to physical ones
- For our purposes, we won't worry about page map size (will store it without using tree structure)

Plan For Today

- **Recap:** Base and bound, multiple segments, and paging
- Page Map Size
- **Demand Paging**

What should we do in our paging design if we run out of physical memory?

Running Out Of Memory

If memory is in high demand, we could fill up all of memory, since a process needs all its pages in memory to run. What should we do in that case?

- Prohibit further program memory requests until some is freed? Not ideal.
- Another idea – what if we kicked out a page and used that page? We could save a page to disk, use the page for new data, and load the old data back in to a physical page later if it's still needed.

We can make physical memory look larger than it is!

Demand Paging

Overall goal: allow programs to run without all their information in memory.

- Keep in memory the information that is being used.
- Keep unused information on disk in *paging file* (also called backing store, or swap space)
- Move information back and forth as needed.
- Locality – most programs spend most of their time using a small fraction of their code and data

Ideally: we have a memory system with the performance of main memory and the cost/capacity of disk!

Demand Paging – 2 Key Questions

1. What is the process for kicking a page out to disk?
2. How do we choose which page to kick out? (next time!)

Demand Paging – 2 Key Questions

1. **What is the process for kicking a page out to disk?**
2. How do we choose which page to kick out? (next time!)

Demand Paging

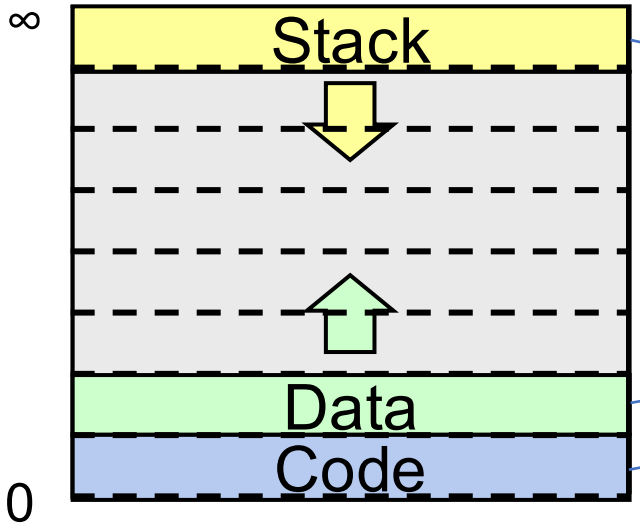

Process A Virtual Address Space

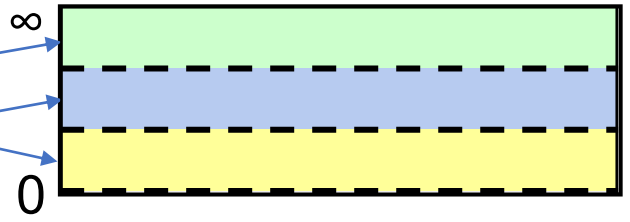

Physical Address Space

	Physical page #	WR?	PR?
7	0	1	1
6	X	X	0
5	X	X	0
4	X	X	0
3	X	X	0
2	X	X	0
1	2	0	1
0	1	0	1

Demand Paging

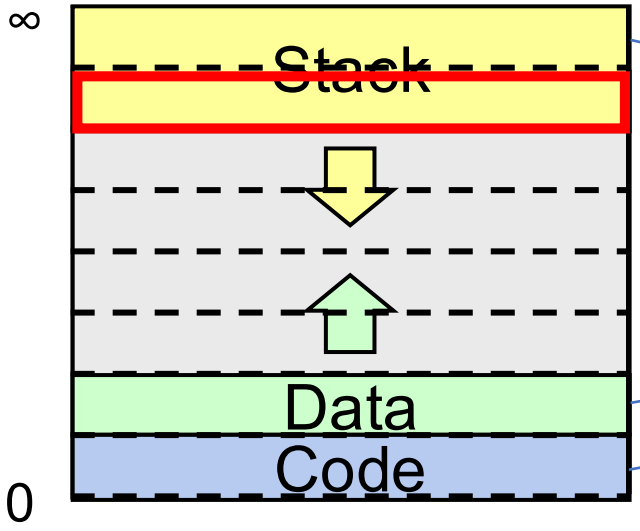

Process A Virtual Address Space

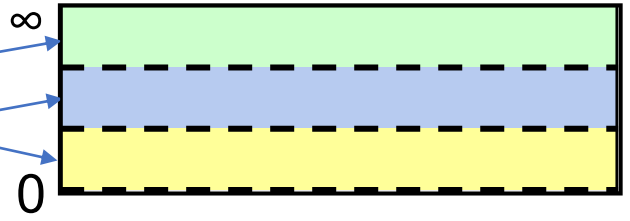

Physical Address Space

	Physical page #	WR?	PR?
7	0	1	1
6	X	X	0
5	X	X	0
4	X	X	0
3	X	X	0
2	X	X	0
1	2	0	1
0	1	0	1

1. Pick an existing physical page and swap it to disk.

Demand Paging

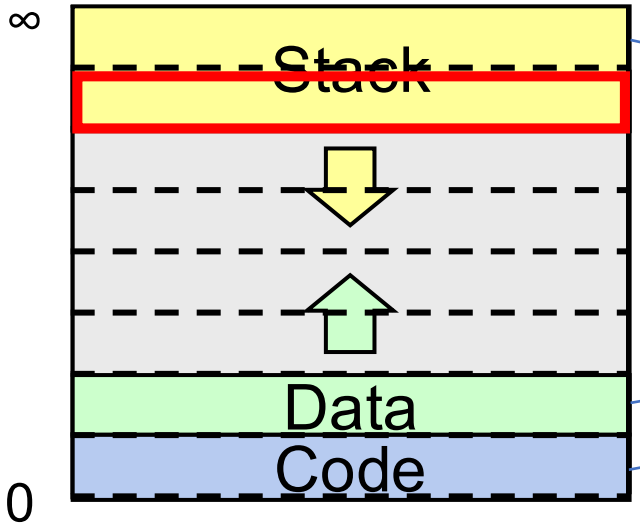

Process A Virtual Address Space

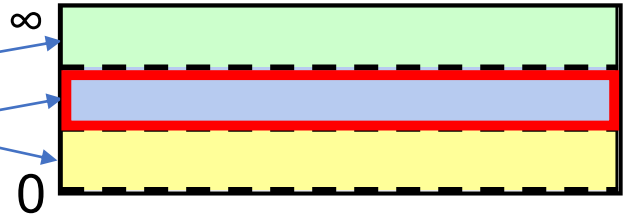

Physical Address Space

	Physical page #	WR?	PR?
7	0	1	1
6	X	X	0
5	X	X	0
4	X	X	0
3	X	X	0
2	X	X	0
1	2	0	1
0	1	0	0

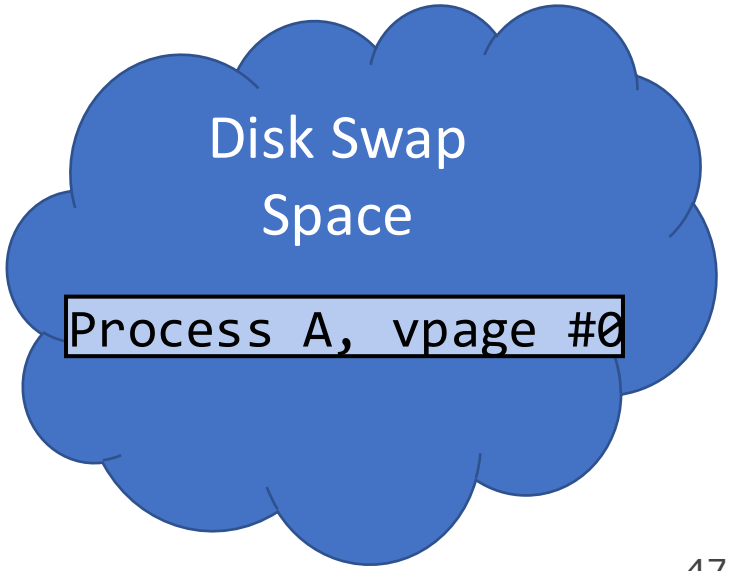

1. Pick an existing physical page and swap it to disk, mark not present.

Demand Paging

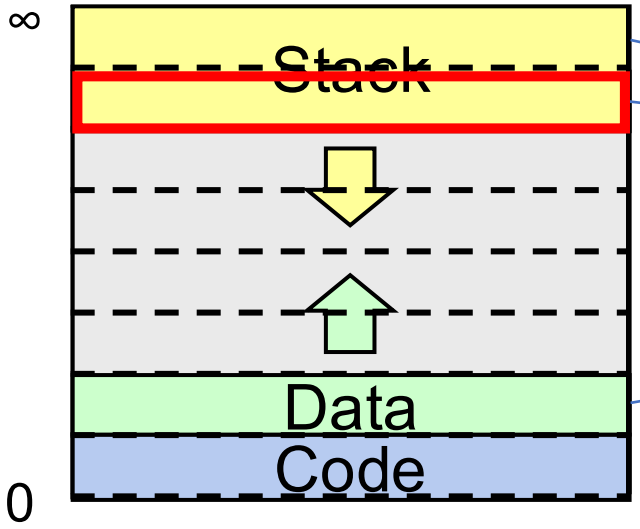

Process A Virtual Address Space

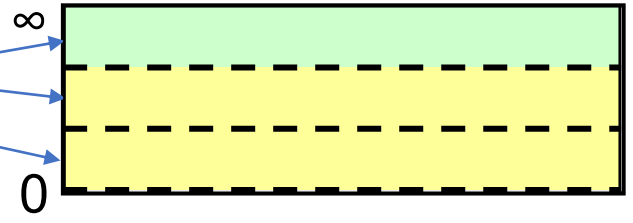

Physical Address Space

	Physical page #	WR?	PR?
7	0	1	1
6	1	1	1
5	X	X	0
4	X	X	0
3	X	X	0
2	X	X	0
1	2	0	1
0	1	0	0

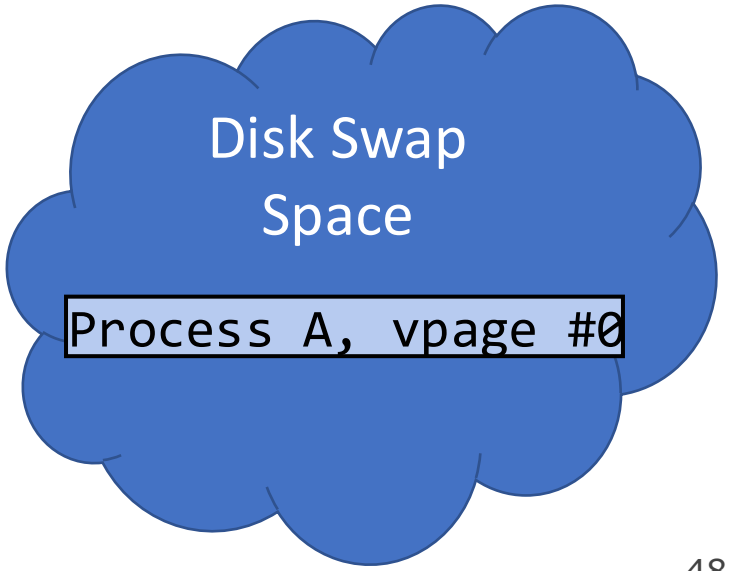

2. Map this physical page to the new virtual page.

Demand Paging

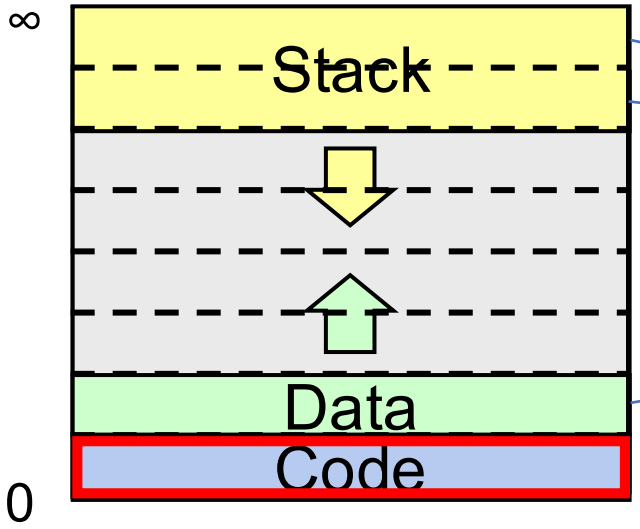

Process A Virtual Address Space

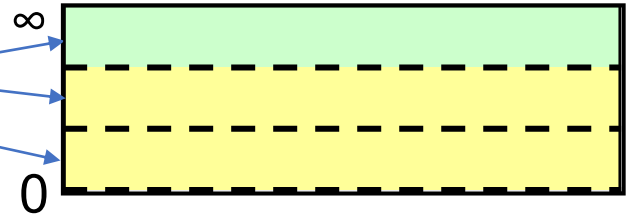

Physical Address Space

	Physical page #	WR?	PR?
7	0	1	1
6	1	1	1
5	X	X	0
4	X	X	0
3	X	X	0
2	X	X	0
1	2	0	1
0	1	0	0

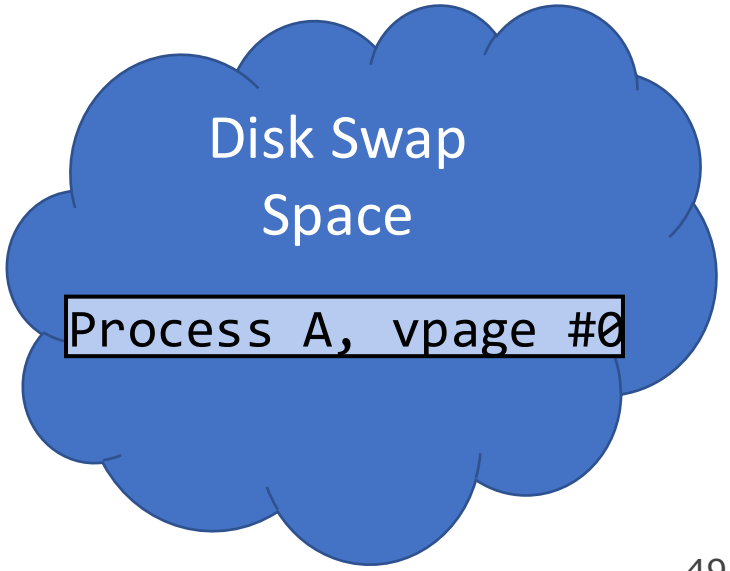

1. We look in the page map and see it's not present.

Demand Paging

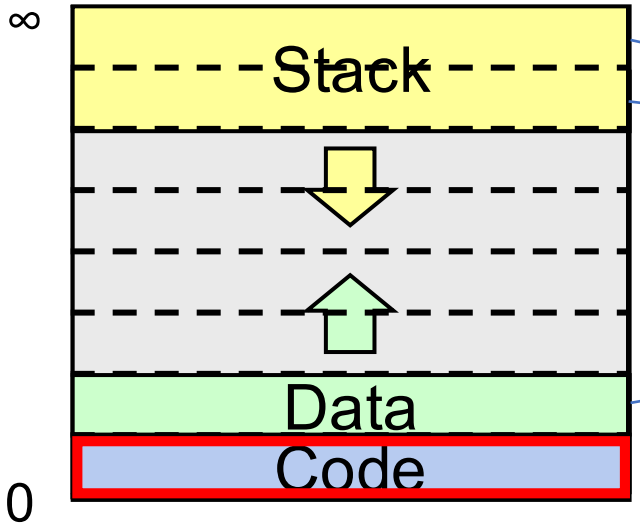

Process A Virtual Address Space

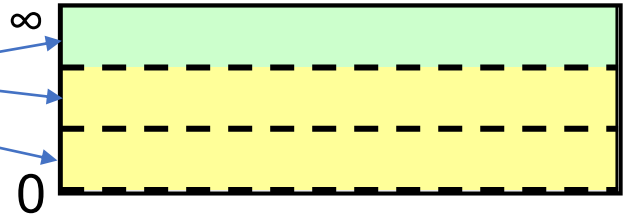

Physical Address Space

	Physical page #	WR?	PR?
7	0	1	1
6	1	1	1
5	X	X	0
4	X	X	0
3	X	X	0
2	X	X	0
1	2	0	1
0	1	0	0

2. But it is stored in disk swap, so we load it back in (kicking another page if needed).

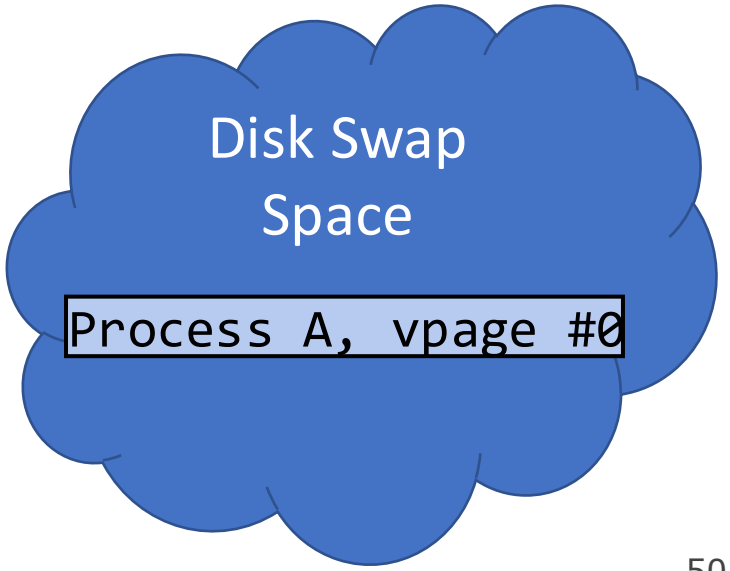

Demand Paging

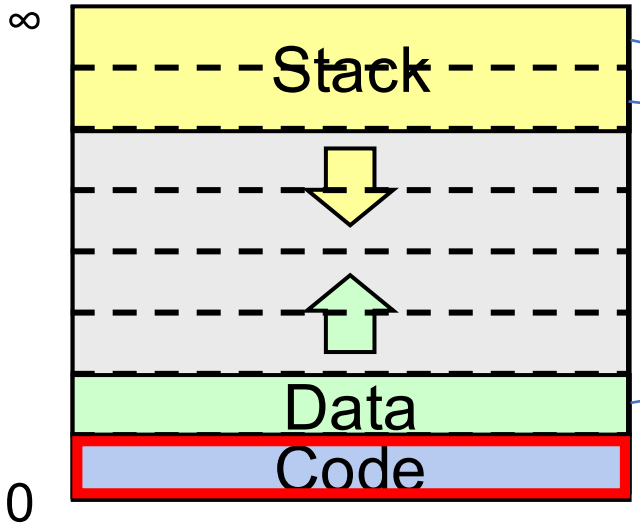

Process A Virtual Address Space

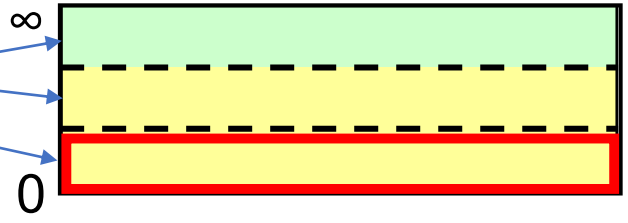

Physical Address Space

	Physical page #	WR?	PR?
7	0	1	0
6	1	1	1
5	X	X	0
4	X	X	0
3	X	X	0
2	X	X	0
1	2	0	1
0	1	0	0

2. But it is stored in disk swap, so we load it back in (kicking another page if needed).

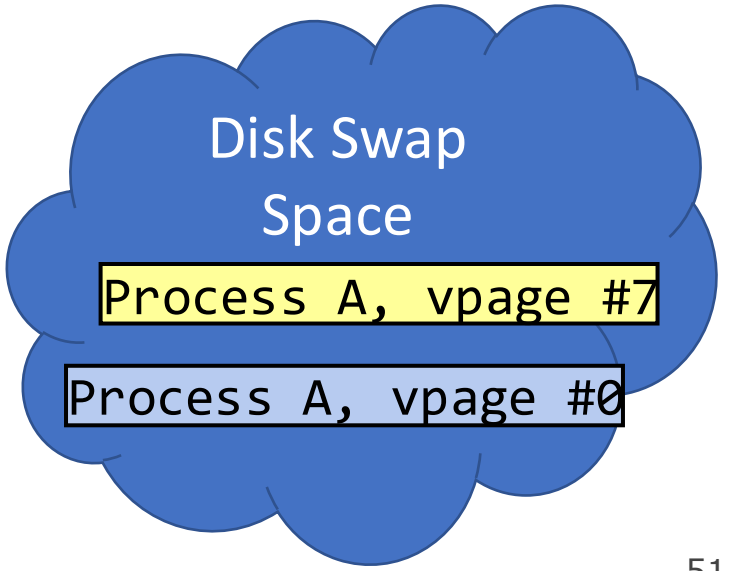

Demand Paging

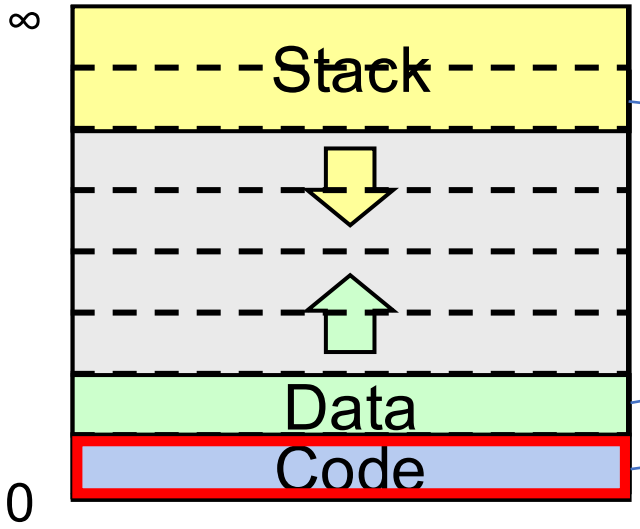

Process A Virtual Address Space

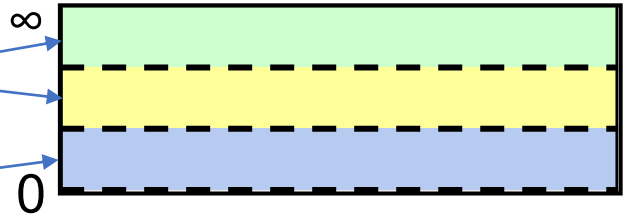

Physical Address Space

	Physical page #	WR?	PR?
7	0	1	0
6	1	1	1
5	X	X	0
4	X	X	0
3	X	X	0
2	X	X	0
1	2	0	1
0	0	0	1

2. But it is stored in disk swap, so we load it back in (kicking another page if needed).

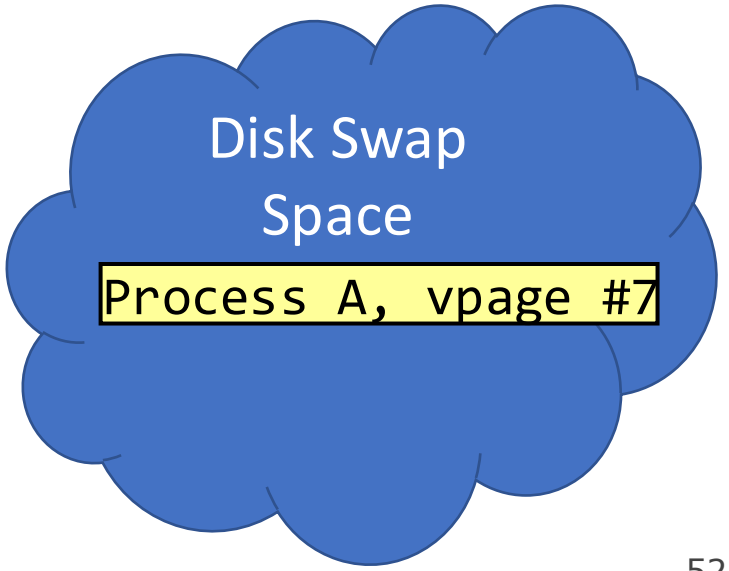

Demand Paging

If we need another page but memory is full:

1. Pick a page to kick out
2. Write it to disk
3. Mark the old page map entry as not present
4. Update the new page map entry to be present and map to this physical page

Demand Paging

If the program accesses a page that was swapped to disk:

1. Triggers a page fault (not-present page accessed)
2. We see disk swap contains data for this page
3. Get a new physical page (perhaps kicking out another one)
4. Load the data from disk into that page
5. Update the page map with this new mapping

We now rely on the present bit to tell us more generally if the page is present in physical memory or not – if not, could still be a valid access if in disk swap.

Thrashing

Demand paging can provide big benefits – but can also run into problematic scenarios. In the following scenario, what is a bad scenario for demand paging for what pages this process accesses next?

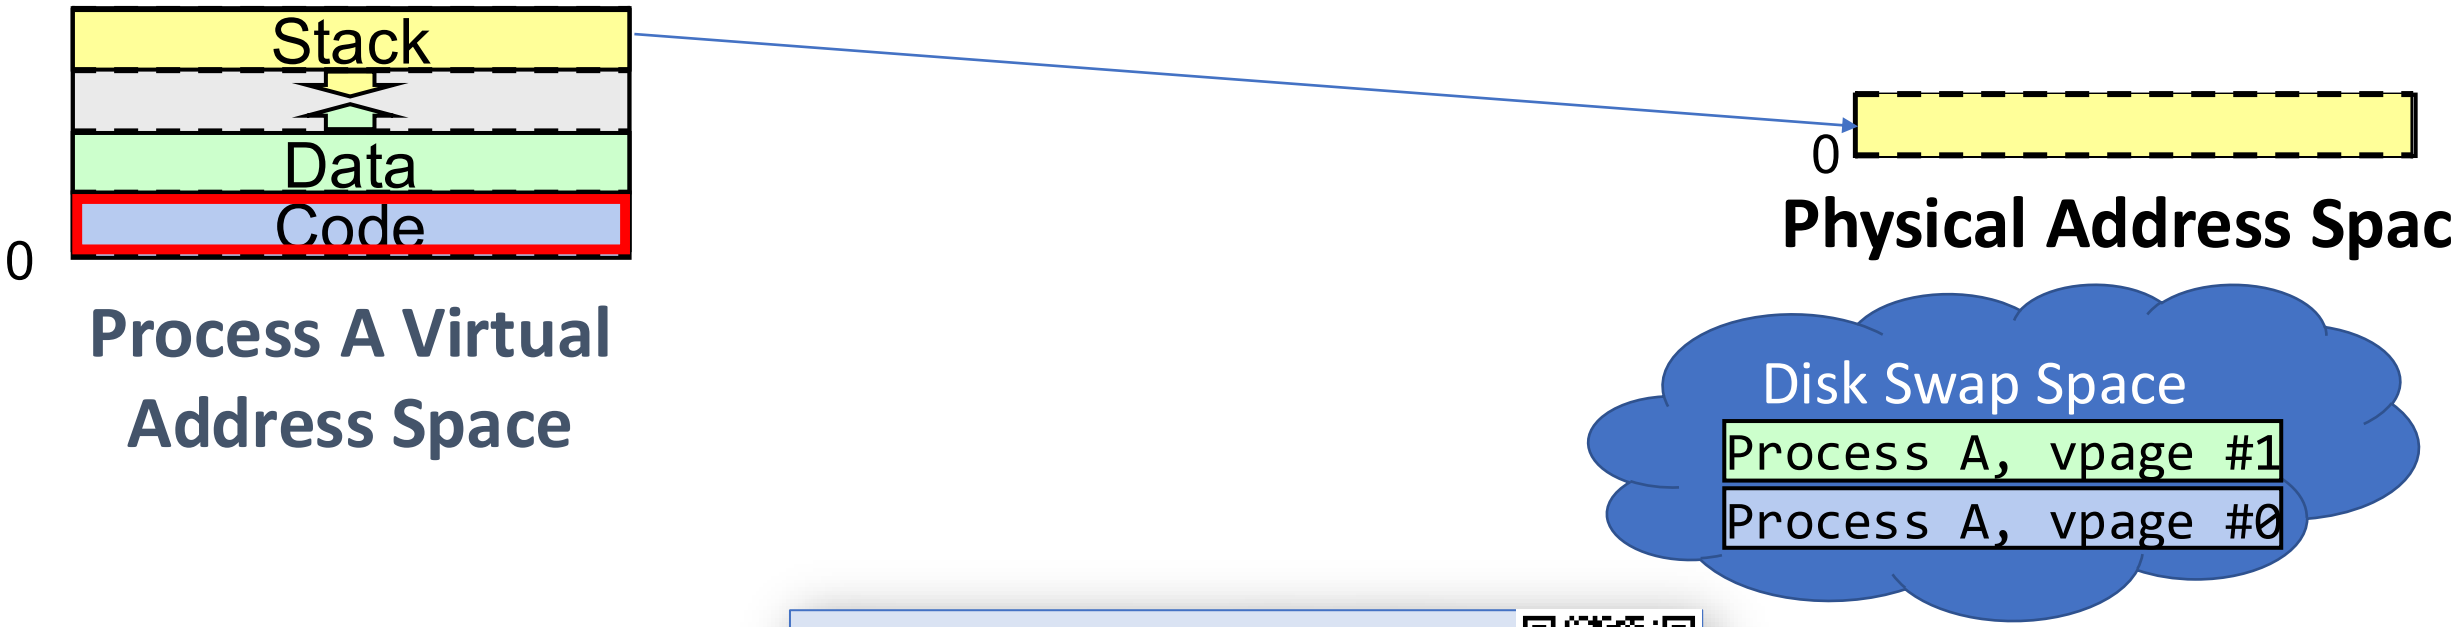

Respond on PollEv:
pollev.com/cs111

L23. Of the following options for what pages this process accesses next, which is the worst for demand paging?

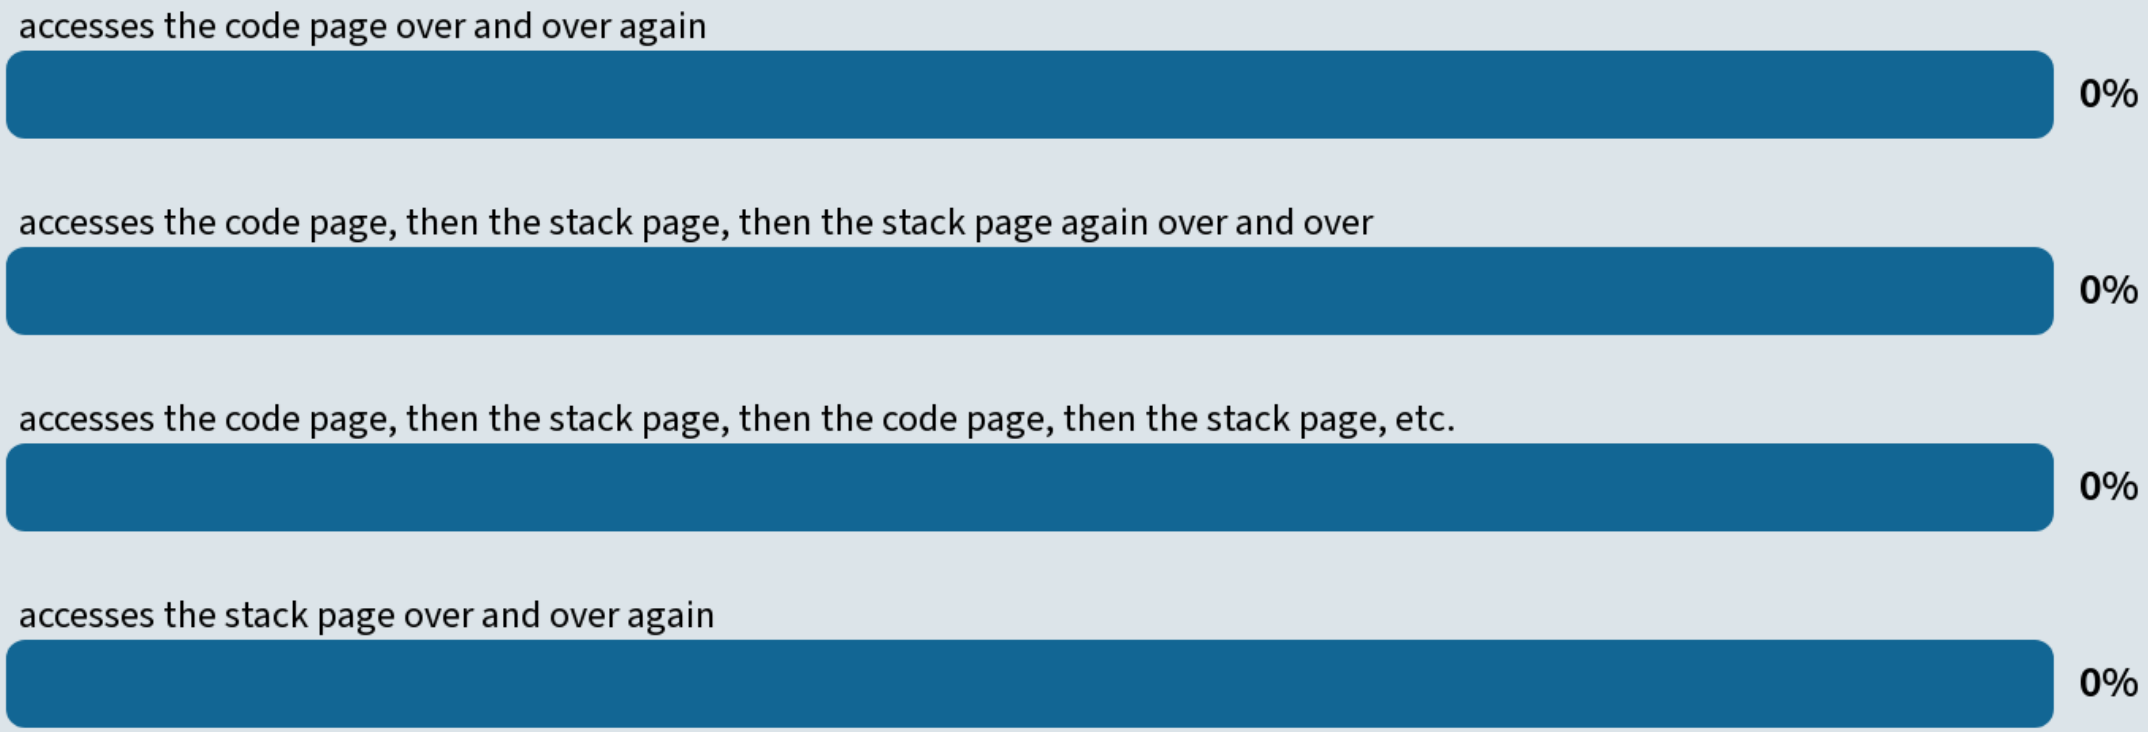

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Thrashing

Demand paging can provide big benefits – but what potential scenario would lead demand paging to slow the system way down?

If the pages being actively used don't all fit in memory, the system will spend all its time reading and writing pages to/from disk and won't get much work done.

- Called *thrashing*
- The page we kick to disk will be needed very soon, so we will bring it back and kick another page, which will be needed very soon, etc….
- Progress of the program will make it look like access time of memory is as slow as disk, rather than disk being as fast as memory. ☹
- With personal computers, users can notice thrashing and kill some processes

Page Fetching

Now we have a mechanism to allow programs to run without all their information in memory. But even if there is space, when should we bring pages into memory?

- Most modern OSes start with no pages loaded, load pages when referenced ("demand fetching").
- Alternative: *prefetching* - try to predict when pages will be needed and load them ahead of time (requires predicting the future…)

Page Replacement

If we need another physical page but all memory is used, **which page should we throw out?**

More next time…

Recap

- **Recap:** Base and bound, multiple segments, and paging
- Page Map Size
- Demand Paging

Next time: how to choose which pages to swap to disk (the clock algorithm).

Lecture 23 takeaway: We can make memory appear larger than it is by swapping pages to disk when we need more space and swapping them back later. But thrashing can occur when the system spends all its time doing disk operations and little time on actual work.