

CS111, Lecture 24

The Clock Algorithm



masks recommended

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)

Announcements

- assign6 released tomorrow
- Friday's lecture: virtualization and trust + assign6 YEAH

Topic 4: 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?

CS111 Topic 4: Virtual Memory



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

Learning Goals

- Learn about tradeoffs in approaches for choosing pages to kick out of memory
- Walk through the implementation of the clock algorithm, one algorithm for choosing which page to throw out

Plan For Today

- **Recap:** Demand Paging
- The Clock Algorithm
- What about when the OS Runs?
- Virtual Memory summary

Plan For Today

- **Recap: Demand Paging**
- The Clock Algorithm
- What about when the OS runs?
- Virtual Memory summary

Demand Paging

If memory is in high demand, we could fill up all of memory, since a process needs all its pages in memory to run.

Thought: does a process really need all its allocated pages in memory?

Let's say memory is full, and a process wants another page. We could “borrow” a used physical page – we'll store its existing contents on disk, and then use the page for this new data. If the old contents are referenced later, we'll load them back into a physical page.

Overall goal: make physical memory look larger than it is.

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

Disk Swap

We don't always need to write a swapped-out page to disk – e.g. read-only code pages can always be loaded from executable. And we may have initial data for a page that wasn't previously swapped out.

There are three categories of pages for swapping to disk:

- 1. Read-only code pages:** read from executable when needed
- 2. Initialized data pages:** on first access, read from executable. Once loaded, save to swap file since contents may have changed.
- 3. Uninitialized data pages:** e.g. stack, heap – on first access, just clear memory to all zeros. Save to swap file as needed.

Disk Swap

We don't always need to write a swapped-out page to disk – e.g. read-only code pages can always be loaded from executable. And we may have initial data for a page that wasn't previously swapped out.

On assign6:

- You'll only write to disk if a page is “dirty” (modified). Page maps contain a dirty bit that is set whenever a page is modified.
- A page may have contents on disk from the executable or from a previous swap – you'll read into memory in both cases.

Plan For Today

- **Recap:** Demand Paging
- **The Clock Algorithm**
- What about when the OS runs?
- Virtual Memory summary

Page Replacement

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

- Random? (works surprisingly well!)
- FIFO? (throw out page that's been in memory the longest) – fairness
- Would be nice if we could pick page whose next access is farthest in the future, but we'd need to predict the future...
- LRU (least-recently-used)? Replace page that was accessed the longest time ago.

Page Replacement

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

- Random? (works surprisingly well!)
- FIFO? (throw out page that's been in memory the longest) – fairness
- Would be nice if we could pick page whose next access is farthest in the future, but we'd need to predict the future...
- **LRU (least-recently-used)? Replace page that was accessed the longest time ago.**

Least-Recently-Used

How could we know which page was the least-recently used?

- Store clock time for each page on each reference?
- Scan all pages to find oldest one?

Alternative: just find an old page, not necessarily the oldest.


The **clock algorithm** is one implementation of this idea.

Clock algorithm key idea: rotate through pages until we find one that hasn't been referenced since the *last time* we checked it. (“second chance algorithm”)

Clock Algorithm

Let's say the program requests mapping page 5, but we have no more physical pages. This triggers the clock algorithm.

The reference bit is set to 1 whenever that page is read or written.




Physical Pages

“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	1
6	D	1	1	1
5	X	X	0	X
4	X	X	0	X
3	X	X	0	X
2	C	1	1	1
1	B	0	1	1
0	A	0	1	1

Page Map

Clock Algorithm



Was this page accessed recently (reference = 1)?
If so, set reference = 0 and continue.

“We’ll leave this page for now – but if we come back and it’s still unused, we’ll kick it out.”


“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	1
6	D	1	1	1
5	X	X	0	X
4	X	X	0	X
3	X	X	0	X
2	C	1	1	1
1	B	0	1	1
0	A	0	1	1

Physical Pages

Page Map

Clock Algorithm



Was this page accessed recently (reference = 1)?
If so, set reference = 0 and continue.

“reference” bit




	Physical page #	WR?	PR?	R
7	E	1	1	1
6	D	1	1	1
5	X	X	0	X
4	X	X	0	X
3	X	X	0	X
2	C	1	1	1
1	B	0	1	1
0	A	0	1	0

Physical Pages

Page Map

Clock Algorithm

Was this page accessed recently (reference = 1)?
If so, set reference = 0 and continue.



Physical Pages


“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	1
6	D	1	1	1
5	X	X	0	X
4	X	X	0	X
3	X	X	0	X
2	C	1	1	1
1	B	0	1	0
0	A	0	1	0

Page Map

Clock Algorithm

Was this page accessed recently (reference = 1)?
If so, set reference = 0 and continue.



Physical Pages


“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	1
6	D	1	1	1
5	X	X	0	X
4	X	X	0	X
3	X	X	0	X
2	C	1	1	0
1	B	0	1	0
0	A	0	1	0

Page Map

Clock Algorithm

Was this page accessed recently (reference = 1)?
If so, set reference = 0 and continue.



Physical Pages

“reference” bit




	Physical page #	WR?	PR?	R
7	E	1	1	1
6	D	1	1	0
5	X	X	0	X
4	X	X	0	X
3	X	X	0	X
2	C	1	1	0
1	B	0	1	0
0	A	0	1	0

Page Map

Clock Algorithm

Was this page accessed recently (reference = 1)?
If so, set reference = 0 and continue.




Physical Pages

“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	0
6	D	1	1	0
5	X	X	0	X
4	X	X	0	X
3	X	X	0	X
2	C	1	1	0
1	B	0	1	0
0	A	0	1	0

Page Map

Clock Algorithm



Was this page accessed recently (reference = 1)?
If not, this is the one we should remove.

“This page hasn’t been used since the last time I checked – let’s remove it.”

“reference” bit




	Physical page #	WR?	PR?	R
7	E	1	1	0
6	D	1	1	0
5	X	X	0	X
4	X	X	0	X
3	X	X	0	X
2	C	1	1	0
1	B	0	1	0
0	A	0	1	0

Page Map

Clock Algorithm

Now the clock algorithm stops, and **we remember the position of the hand for next time it runs.**



Physical Pages


“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	0
6	D	1	1	0
5	A	1	1	1
4	X	X	0	X
3	X	X	0	X
2	C	1	1	0
1	B	0	1	0
0	A	0	0	0

Page Map

Clock Algorithm

Let's say the program now requests mapping page 4. Some memory accesses have also happened.



Physical Pages


“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	0
6	D	1	1	0
5	A	1	1	1
4	X	X	0	X
3	X	X	0	X
2	C	1	1	1
1	B	0	1	1
0	A	0	0	0

Page Map

Clock Algorithm

Was this page accessed recently (reference = 1)?
If so, set reference = 0 and continue.



Physical Pages


“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	0
6	D	1	1	0
5	A	1	1	1
4	X	X	0	X
3	X	X	0	X
2	C	1	1	1
1	B	0	1	0
0	A	0	0	0

Page Map

Clock Algorithm

Was this page accessed recently (reference = 1)?
If so, set reference = 0 and continue.



Physical Pages

“reference” bit




	Physical page #	WR?	PR?	R
7	E	1	1	0
6	D	1	1	0
5	A	1	1	1
4	X	X	0	X
3	X	X	0	X
2	C	1	1	0
1	B	0	1	0
0	A	0	0	0

Page Map

Clock Algorithm

Was this page accessed recently (reference = 1)?
If not, this is the one we should remove.

“This page hasn’t been used since the last time I checked – let’s remove it.”



Physical Pages


“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	0
6	D	1	1	0
5	A	1	1	1
4	X	X	0	X
3	X	X	0	X
2	C	1	1	0
1	B	0	1	0
0	A	0	0	0

Page Map

Clock Algorithm

Now the clock algorithm stops, and **we remember the position of the hand for next time it runs.**



Physical Pages

“reference” bit


	Physical page #	WR?	PR?	R
7	E	1	1	0
6	D	1	0	0
5	A	1	1	1
4	D	1	1	1
3	X	X	0	X
2	C	1	1	0
1	B	0	1	0
0	A	0	0	0

Page Map

Clock Algorithm

- We add a *reference* bit: set whenever a page is read or written
- When physical memory is full and we need to choose a page to remove, run the clock algorithm.
- Clock hand “sweeps” over pages, rotating back to start if reaching the end.
- Every time the hand visits a page, we ask: “Has this page been referenced since the last time the clock hand swept over it?”
 - **If YES (reference = 1):** mark it as not referenced, and advance clock hand
 - **If NO (reference = 0):** choose it for removal, advance clock hand, stop clock algorithm
- The clock hand position is saved for the next time the algorithm runs
- “Second chance” algorithm

Clock Algorithm



Some time has passed, pages were referenced, and we now need a new page. Which page will the clock algorithm choose to reuse this time? (hand starts at E)

“reference” bit




	Physical page #	WR?	PR?	R
7	E	1	1	1
6	D	1	0	0
5	A	1	1	1
4	D	1	1	1
3	A	1	0	0
2	C	1	1	0
1	B	0	1	1
0	A	0	0	0

Physical Pages

Page Map

Clock Algorithm



Some time has passed, pages were referenced, and we now need a new page. Which page will the clock algorithm choose to reuse this time? (hand starts at E)

“reference” bit

	Physical page #	WR?	PR?	R
7	E	1	1	1
6	D	1	0	0
5	A	1	1	1
4	D	1	1	1
3	A	1	0	0
2	C	1	1	0
1	B	0	1	1
0	A	0	0	0

Physical Pages

Page Map

Page Replacement

How does page replacement work if there are multiple processes running?

- *Per-process replacement*: each process has separate pool of physical pages, and a page fault in a process can only replace one of its own pages. But how many physical pages should each process get?
- *Global replacement* (most common): all pages from all processes in single replacement pool. A page fault in one process can kick out a page in another process.

Plan For Today

- Recap: Demand Paging
- The Clock Algorithm
- **What about when the OS runs?**
- Virtual Memory summary

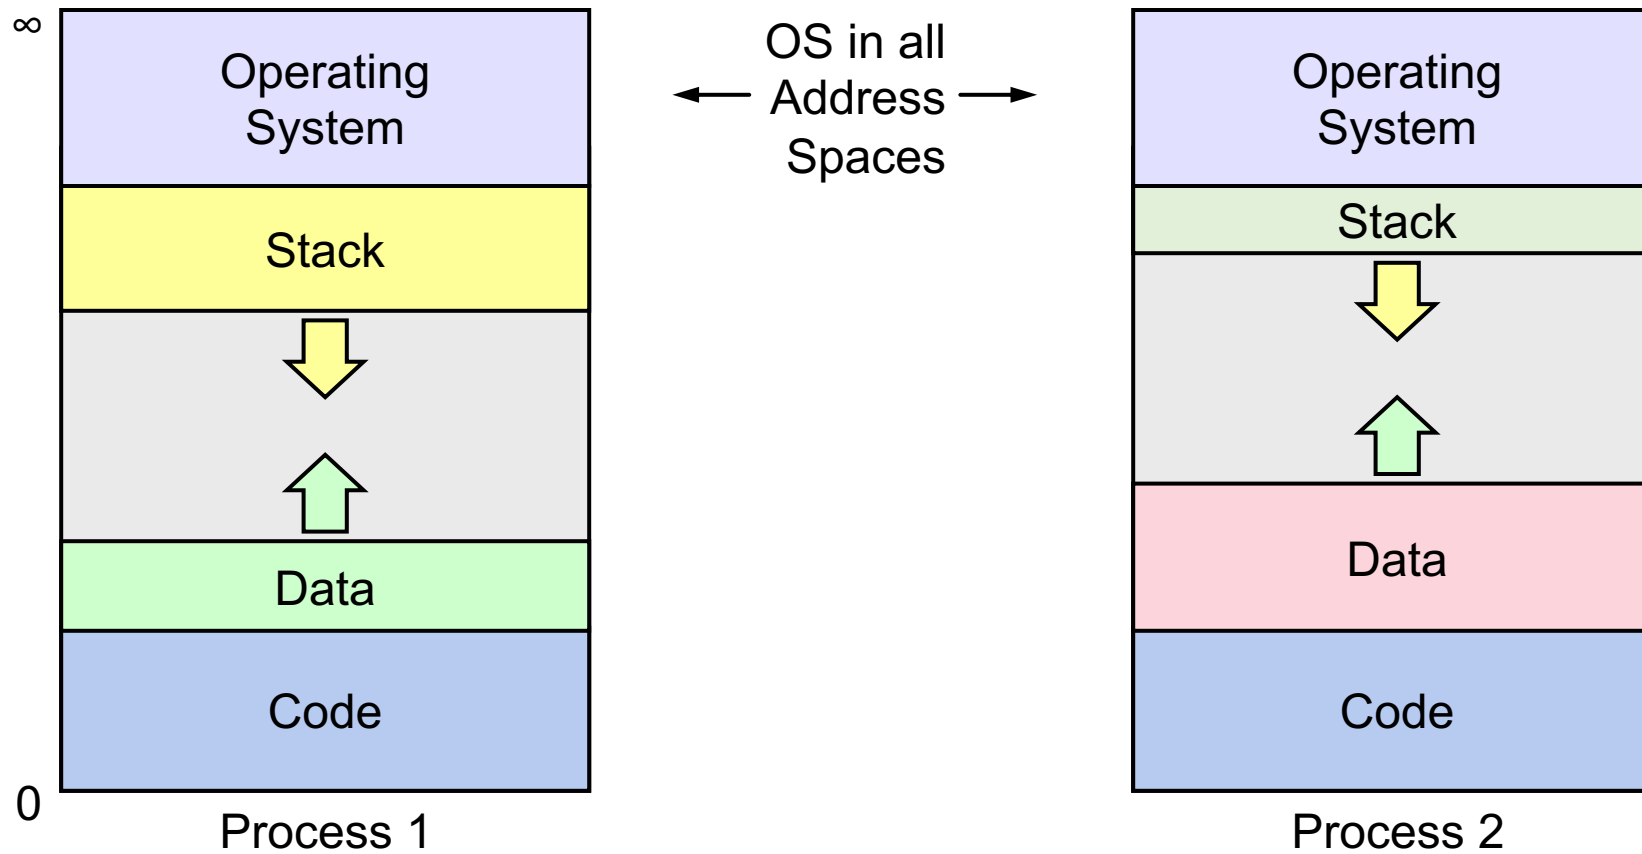
OS Execution

How does virtual memory work when the OS runs?

Challenge: whenever we make a system call, we pass *virtual* addresses from our process, or the OS may need to access data from the process's virtual address space.

Solution on most modern systems: OS has space in every process's virtual address space.

OS and User in Same Address Space



OS Execution

How does virtual memory work when the OS runs?

OS has space in every process's virtual address space.

Problem: don't want user program accessing OS pages.

Solution: new bit in page table that marks kernel-only pages. When in user mode, not accessible, but accessible when OS is running.

Plan For Today

- **Recap:** Demand Paging
- The Clock Algorithm
- What about when the OS runs?
- **Virtual Memory summary**

Virtual Memory

- Virtual memory is an example of “OS magic” – very powerful mechanism
- Virtualization: making one thing look like another – separation between appearance and reality
- OS can manage physical memory how it wants (e.g. swap to disk), invisible to user programs

Goals:

- **Multitasking** – allow multiple processes to be memory-resident at once
- **Transparency** – no process should need to know memory is shared. Each must run regardless of the number and/or locations of processes in memory.
- **Isolation** – processes must not be able to corrupt each other
- **Efficiency** (both of CPU and memory) – shouldn't be degraded badly by sharing

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?*

Why is answering this question important?

- We can understand one of the most “magical” responsibilities of OSes – making one set of memory appear as several!
- Exposes challenges of allowing multiple processes share memory while remaining isolated
- Allows us to understand exactly what happens when a program accesses a memory address

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

Recap

- Recap: Demand Paging
- The Clock Algorithm
- What about when the OS runs?
- Virtual Memory summary

Lecture 24 takeaway: There are many different policies to choose a page to kick out when memory is full. The clock algorithm is one approximation of LRU to pick an old page to remove.