# Virtual Memory (III) — Policies

Hung-Wei Tseng



### **Recap: Demand paging**





### data 0x80008000

### Program

Ծ

 $\mathcal{O}$ nstructi

0f00bb27 509cbd23 00005d24 0000bd24 2ca422a0 130020e4 00003d24 2ca4e2b3

00c2e800 00000008 00c2f000 80000008 00c2f800 00000008 00c30000 00000008



### **Recap: Hierarchical page table to make paging feasible**



- Make page table "nodes" demand-pagable -1 not all of them has to
- root node of the hierarchical page table

# **Recap: Page replacement policy**

- Goal: Identify page to remove that will avoid future page faults (i.e. utilize) locality as much as possible)
- Implementation Goal: Minimize the amount of software and hardware overhead
  - Example:
    - Memory (i.e. RAM) access time: 100ns
    - Disk access time: 10ms
    - P<sub>f</sub>: probability of a page fault
    - Effective Access Time =  $10^{-7} + P_f * 10^{-3}$
  - When  $P_f = 0.001$ : Effective Access Time = 10,100ns
  - Takeaway: Disk access tolerable only when it is extremely rare





### What VAX/VMS proposed to achieve these goals?

 Considering the optimization goals and the proposed VAX/ VMS mechanisms, which of the following combinations is incorrect?

	Goal		Opt
	Process startup cost	W	Demand-zero
P	Process performance interference	Χ	Process-local
С	Page table lookup overhead	Y	Page clusterin
D	Paging load on disks	Ζ	Page caching

- We're still using their proposed techniques almost everyday!
- It's basically the baseline UNIX VM design

### timization

- & copy-on-refernce
- replacement
- also helps reduce disk loads

### Machine-Independent Virtual Memory Management for Paged Uniprocessor and Multiprocessor Architectures Richard Rashid, Avadis Tevanian, Michael Young, David Golub, Robert Baron, David Black,

Richard Rashid, Avadis Tevanian, Michael Young, David Golub, Robert Bard William Bolosky, and Jonathan Chew Carnegie-Mellon University, NeXT, University of Rochester

### **Recap: Overview of Mach's VM**





a memory object could be anything — a file, a network buffer, remote network memory, device buffer, or physical DRAM

### **Recap: Address allocation is sparse in multithreading model!**



address map



### Outline

- Page replacement policies
- Page replacement policy once used in UNIX: Converting a Swap-Based System to do Paging in an Architecture Lacking Page-Reference Bits
- Another popular page replacement policy: WSClcok A Simple and Effective Algorithm for Virtual Memory Management

# Page replacement policies from textbooks

# Page replacement policy

- We need to determine:
  - Which page(s) to remove
  - When to remove the page(s)
- Goals
  - Identify page to remove that will avoid future page faults (i.e. utilize) locality as much as possible)
  - Minimize the amount of software and hardware overhead



# **Page replacement algorithms**

- FIFO: Replace the oldest page
- LRU: Replace page that was the least recently used (longest) since last use)



### https://www.pollev.com/hungweitseng close in 1:00 **FIFO v.s. LRU**

 Assume your OS uses FIFO policy when handle page faults. Also assume that we have 3 physical memory pages available. Compared with the same machine using an OS with LRU page replacement police, how many more page faults will you see for the FIFO based OS in the following page reference sequence?



	10	11
>	5	2





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



### https://www.pollev.com/hungweitseng close in 1:00 **FIFO v.s. LRU**

 Assume your OS uses FIFO policy when handle page faults. Also assume that we have 3 physical memory pages available. Compared with the same machine using an OS with LRU page replacement police, how many more page faults will you see for the FIFO based OS in the following page reference sequence?

			0	1	2	3	4	5	6	7	8	
		Page #	2	3	2	1	5	2	4	5	3	
A.	0										1	
B.	1											A
D.	2											E
E.	4											E



iotal Results: (

# FIFO (Group)



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



# **FIFO v.s. LRU**

 Assume your OS uses FIFO policy when handle page faults. Also assume that we have 3 physical memory pages available. Compared with the same machine using an OS with LRU page replacement police, how many more page faults will you see for the FIFO based OS in the following page reference sequence?



# **FIFO v.s. LRU**





May require hardware support or linked list or additional timestamps in page tables

High — you need to manipulate the list or update every counter

Usually better than FIFO

# Converting a Swap-Based System to do Paging in an Architecture Lacking Page-Reference Bits

Özalp Babaoglu and William Joy\* Cornell University and University of California, Berkeley

### https://www.pollev.com/hungweitseng close in 1:00 **The VMS/Old UNIX VM**

- Regarding the original UNIX VM (basically the VMS), please identify how many of the following statements are correct.
  - ① VAX machine provides no hardware support for page replacement policies
  - VMS implements FIFO policy for page replacement 2
  - ③ A process's resident set cannot be adjusted even though that process is the only process in the system
  - VMS swaps out all memory page belong to a process when that process is (4)switched out





VMS	

atal Results:







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



### https://www.pollev.com/hungweitseng close in 1:00 **The VMS/Old UNIX VM**

- Regarding the original UNIX VM (basically the VMS), please identify how many of the following statements are correct.
  - ① VAX machine provides no hardware support for page replacement policies
  - 2 VMS implements FIFO policy for page replacement
  - ③ A process's resident set cannot be adjusted even though that process is the only process in the system
  - VMS swaps out all memory page belong to a process when that process is (4)switched out





VMS (Group)					

stal Results:

# VMS (Group)





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



# The VMS/Old UNIX VM

- Regarding the original UNIX VM (basically the VMS), please identify how many of the following statements are correct.
  - ① VAX machine provides no hardware support for page replacement policies
  - ② VMS implements FIFO policy for page replacement
  - ③ A process's resident set cannot be adjusted even though that process is the only process in the system
  - VMS swaps out all memory page belong to a process when that process is Really inefficient if you have frequent context switches or if you have many applications in-fly switched out Whenever a process is removed from memory, its entire (4)
  - resident set is written to the swap file, along with some Α. ()their decisions on. Without even this tions B. minimal page reference information, the only reasonable algorithms for replacing pages are the C. 2 First-In-First-Cut (FIFO) and the Random (RAND) D. 3



In VMS, the vendor-supplied operating system for the VAX, the solution to the replacement decision is simple. Each process is assigned a fixedsize memory partition, called a resident set, that is managed according to the FIFU DOILCY. rages

# The Why of Babaoglu new UNIX VM

- The original UNIX is a swap-based system
  - Whenever you have a context switch, swap the whole process out from the memory
  - Really inefficient if you have frequent context switches or if you have many applications in-fly
  - Imply that the modern UNIX or Linux does not do this
- Efficient page replacement policies and other virtual optimization techniques cannot be implemented easily without appropriate hardware support



### https://www.pollev.com/hungweitseng close in 1:00

# The page replacement policy proposed

- How many of following statements fit the page replacement policy that the paper implements?
  - ① It uses LRU (least-recently-used) as the page replacement policy
  - ② Page replacement policy are only triggered whenever a page fault occurs
  - ③ It attaches a timestamp to each page table entry instead of using the reference bit from hardware
  - ④ Processes are allocated a fixed set of pages and swap in/out to/from those pages
  - ⑤ The page replacement policy helps to guarantee the response time of short programs
  - A. 0
    B. 1
    C. 2
    D. 3
    E. 4



- olicy fault occurs sing the reference bit
- to/from those pages e time of short programs

### UNIXVM

### UNIXVM





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



### https://www.pollev.com/hungweitseng close in 1:00

# The page replacement policy proposed

- How many of following statements fit the page replacement policy that the paper implements?
  - ① It uses LRU (least-recently-used) as the page replacement policy
  - ② Page replacement policy are only triggered whenever a page fault occurs
  - ③ It attaches a timestamp to each page table entry instead of using the reference bit from hardware
  - ④ Processes are allocated a fixed set of pages and swap in/out to/from those pages
  - ⑤ The page replacement policy helps to guarantee the response time of short programs
  - A. 0 B. 1 C. 2 D. 3 E. 4

- olicy fault occurs sing the reference bit
- to/from those pages e time of short programs

### UNIXVM (Group)

# UNIXVM (Group)



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





G

Ε

attach a "reference bit" to each PTE, set to true when the page is referenced



С







С









С



Α

Where to put M ?



**Clock hand move** sequentially to swap out the first page without reference bit set. Clear the reference bit when it's set



L

G

BR

F











### C will be selected to swap out, but Rs of A and B are cleared





# **Recap: LRU**

 Assume your OS uses LRU policy when handle page faults. Also assume that we have 3 physical memory pages available. How many page faults will you see in the following page reference sequence?



	9	10	11
	2	5	2
	3	3	
	5	5	
)	2	2	
### https://www.pollev.com/hungweitseng close in 1:00

## How good is clock?

 Assume your OS uses the clock policy when handle page faults. Also assume that we have 3 physical memory pages available. How many page faults will you see in the following page reference sequence?

		Ο	1	2	3	4	5	6	7	8	9	10	11
	Page #	2	3	2	1	5	2	4	5	3	2	5	2
A. 5													
B. 6									1				С
C. 7										A	١.		
D. 8										E	3		
E. 9											, )		
										E			

Clock	





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



### Total Results: 0

### https://www.pollev.com/hungweitseng close in 1:00

## How good is clock?

 Assume your OS uses the clock policy when handle page faults. Also assume that we have 3 physical memory pages available. How many page faults will you see in the following page reference sequence?

		0	1	2	3	4	5	6	7	8	9	10	11
	Page #	2	3	2	1	5	2	4	5	3	2	5	2
A. 5													
B. 6									1			(	Clocl
C. 7										A	ι		
D. 8										E	3		
E. 9											; )		
										E			

### Clock (Group)



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



### Total Results: 0

## How good is clock?

 Assume your OS uses the clock policy when handle page faults. Also assume that we have 3 physical memory pages available. How many page faults will you see in the following page reference sequence?



9	10	11
2	5	2
<b>0</b> .	2+.	
Z* 5+	Z** 5+	
3	3	

### **Recap: Page caching to cover the performance loss**

- Evicted pages will be put into one of the lists in DRAM
  - Free list: clean pages
  - Modified list: dirty pages needs to copy data to the disk
- Page fault to any of the page in the lists will bring the page back
  - Reduces the demand of accessing disks



## Free list in Babaoglu's UNIX

- How many of the following statements regarding the "free list" is/are correct?
  - ① It can improve the latency of a page fault
  - It can reduce the latency of swapping out a page
  - ③ It can incur disk accesses without page faults
  - ④ It doesn't allow a page in the list to be used for other purpose
  - A. 0
  - B. 1
  - C. 2
  - D. 3 E. 4



https://www.pollev.com/hungweitseng close in 1:00

### Freelist

idal Results:

Powered by D Poll Everywhere . Get help at policy.com/spp Suit the presentation to







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



### Total Results: 0

## Free list in Babaoglu's UNIX

- How many of the following statements regarding the "free list" is/are correct?
  - ① It can improve the latency of a page fault
  - It can reduce the latency of swapping out a page
  - ③ It can incur disk accesses without page faults
  - ④ It doesn't allow a page in the list to be used for other purpose
  - A. 0
  - B. 1
  - C. 2
  - D. 3 E. 4



https://www.pollev.com/hungweitseng close in 1:00

### Freelist (Group)

idal Results:

Powered by Roll Everywhere n. Get help at policy.com/spp Suit the presentation to

### Freelist (Group)



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



### Total Results: 0

## Free list in Babaoglu's UNIX

- How many of the following statements regarding the "free list" is/are correct?
  - ① It can improve the latency of a page fault instead of swapping a page during the page fault, just take one from the free list
  - 2 It can reduce the latency of swapping out a page No! This completely depend on how fast your disk/storage is!
  - ③ It can incur disk accesses without page faults. Do you remember how UNIX page replacement is triggered?
  - ④ It doesn't allow a page in the list to be used for other purpose
  - A. 0
  - B. 1

  - D. 3



- No! You can use those pages as disk caches!

### **Free list**

- So far, we need to trigger clock policy and swap in/out on each page fault
- Why don't we prepare more free pages each time so that we can feed page faults with pages from the list?
- Free list
  - When we need a page, take one from the free list
  - Have a daemon running the background, managing this free list you can do this when system is not loaded
  - If size of free list gets too small, trigger the clock algorithm to add pages into the free list (by swapping out to disk)
  - Free list can be used as a disk cache

## The page replacement policy proposed

- How many of following statements fit the page replacement policy that the paper implements?
  - ① It uses LRU (least recently used) as the page replacement policy clock
  - ② Page replacement policy are only triggered whenever a page fault occurs
  - It attaches a timestamp to each page table entry instead of using the reference bit from hardware reference bit
  - ④ Processes are allocated a fixed set of pages and swap in/out to/from those pages
  - The page replacement policy helps to guarantee the response time of short programs
  - A. 0
  - B. 1
  - C. 2
  - D. 3
  - E. 4

- olicy <del>fault occurs</del> r a threshold sing the reference bit
- to/from those pages e time of short programs

## The page replacement policy proposed

- How many of following statements fit the page replacement policy that the paper implements?
  - ① It uses LRU (least recently used) as the page replacement policy
  - ② Page replacement policy are only triggered whenever a page fault occurs...
  - ③ It attaches a timestamp to each page table entry instead of using the reference bit from hardware<sup>reference bit</sup>
  - Processes are allocated a fixed set of pages and swap in/out to/from those pages (4)
  - The page replacement policy helps to guarantee the response time of short programs (5)
  - Α. B. 1 C. 2 D. 3 E. 4

formance implications. Lazowska [LAZO 79] reports that in his measurements based on a real workload, system performance was significantly improved by increasing the minimum size of the free list (a system generation parameter). unfortunate An

> The projected workload for the system had no (ii) requirement of guaranteed response times as in real-time applications.

### How to implement a simple shell?

- Say, we want to implement a shell that interprets command line commands and executes "./a"
- The following program can serve for this purpose:

```
int main(int argc, char *argv[]) {
 int child pid;
 char cmd[1024];
 memset(cmd, 0 , 1024);
 fprintf(stderr, "CSC501-myshell$ ");
 while(fgets_wrapper(cmd, 1024, stdin)) {
      if(strcmp("exit", cmd)==0)
          exit(1);
      child_pid = fork();
      if (child_pid == 0)
          execvp(cmd,NULL);
      else {
          fprintf(stderr, "CSC501-myshell$ ");
          memset(cmd, 0 , 1024);
 return 0:
```

• Do we actually need the code segment of the parent?



### What happens on a fork?

- Create a new page table
- Copy data page-by-page
- 80% of fork occurs in shell command interpreter



## virtual fork — vfork

- How many of the following statements regarding "vfork" is/are correct?
  - ① vfork can improve the response time of the parent process
  - ② vfork does not create a new address space upon the creation of new process
  - ③ vfork allows the child process to execute within the parent's memory address
  - ④ The child process created by vfork can potentially corrupt the parent process
  - A. 0
  - B. 1
  - C. 2
  - D. 3

### E. 4

rk" is/are correct? cess reation of new process

## virtual fork — vfork

- How many of the following statements regarding "vfork" is/are correct?
  - ① vfork can improve the response time of the parent process
  - ② vfork does not create a new address space upon the creation of new process
  - ③ vfork allows the child process to execute within the parent's memory address
  - ④ The child process created by vfork can potentially corrupt the parent process
  - A. 0
  - B. 1
  - C. 2
  - D. 3

### E. 4

rk" is/are correct? cess reation of new process

## virtual fork — vfork

- How many of the following statements regarding "vfork" is/are correct?
  - vfork can improve the response time of the parent process
    wfork does not create a new address space upon the creation of new process

    - ③ vfork allows the child process to execute within the parent's memory address
    - The child process created by vfork can potentially corrupt the parent process (4)
    - A. 0
    - B. 1
  - C. 2
  - E. 4

### WSClcok - A Simple and Effective Algorithm for Virtual Memory Management Richard Carr and John Hennessy

### **Brief recap: what policies are used?**

- Local: select one page from the same process' physical pages for storing the demanding page when swapping is necessary
  - VAX/VMS
  - Original UNIX
- Global: select any page that was previously belong to any process when swapping is necessary
  - UNIX after Babaoglu
  - Mach



## **Degree of parallelism and performance**

accomplish most tasks-- processes are cheap.

- How many of the following would happen in Babaoglu's UNIX VM if we keep increase the amount of concurrent processes and assume each process uses some virtual memory in the system?
  - The CPU utilization will keep increasing and stay at 100% (1)
  - The system may spend more time in context switching than real computation (2)
  - The system may spend more time in swap in/out than real computation 3
  - Some process may not respond due to the high paging overhead (4)
  - A. 0
  - B. 1
  - C. 2
  - D. 3
  - E. 4



### https://www.pollev.com/hungweitseng close in 1:00

# These

**Degree of Parallelism** 

otal Results: 0

### **Degree of Parallelism**



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





Total Results: 0

### https://www.pollev.com/hungweitseng close in 1:00 **Degree of parallelism and performance**

- How many of the following would happen in Babaoglu's UNIX VM if we keep increase the amount of concurrent processes and assume each process uses some virtual memory in the system?
  - The CPU utilization will keep increasing and stay at 100% (1)
  - The system may spend more time in context switching than real computation (2)
  - The system may spend more time in swap in/out than real computation 3
  - Some process may not respond due to the high paging overhead (4)
  - A. 0
  - B. 1
  - C. 2
  - D. 3
  - E. 4



Degree of Parallelism (Group)

otal Results:

### **Degree of Parallelism (Group)**



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





Total Results: 0

## **Thrashing: Paging overhead**

- The system overcommitted memory to tasks
- The system spends most time in paging, instead of making meaningful progress

Previously, we have seen how scheduling policies can help improving "saturation". Now, let's see how page replacement policies can address this "thrashing"





### degree of multiprogramming

### Saturation: Context Switch Overhead

### You think round robin should act like this —



### But the fact is —

		P1	Overho P1->	ead P2	P2	Overhea P2 -> P	ad '3	P3	Overhead P3 -> P1	I	P1	Overhead P1 -> P2		P2
0	1		1	2		2	3		3	4		4	5	

- Your processor utilization can be very low if you switch frequently
- •No process can make sufficient amount of progress within a given period of time
- It also takes a while to reach your turn



**Overhead** P2 -> P3

## **Thrashing v.s. Saturation**

- Thrashing when memory are overcommitted
  - The system is busy paging
  - The processor is idle waiting
- Saturation when processors are overcommitted
  - The system is busy context switching and scheduling
  - The processor is busy but not contributing to the running program



## **Degree of parallelism and performance**

- How many of the following would happen in Babaoglu's UNIX VM if we keep increase the amount of concurrent processes and assume each process uses some virtual memory in the system?
  - X The CPU utilization will keep increasing and stay at 100%
  - The system may spend more time in context switching than real computation 2
  - The system may spend more time in swap in/out than real computation 3
  - Some process may not respond due to the high paging overhead (4)
  - A. 0
  - **B**. 1
  - C. 2



## **Why WS-Clock**

- Take advantages from both local and global page replacement policies
  - Global simplicity, adaptive to process demands
  - Local prevent thrashing

## **Working Set Algorithm**

- Working set: the set of pages used in a certain number of recent accesses
- Assume these recently referenced pages are likely to be referenced again soon (temporal locality)
- Evict pages that are not referenced in a certain period of time
  - Swap out may occur even if there is no page faults
- A process is allowed to be executed only if the working set size fits in the physical memory



### **WSClock**

- Use working set policy to decide how many pages can a process use
  - Return a page to the free list if there exists a page in the process' working set that hasn't been access for a certain period of time
- If the free list is lower than a threshold
  - Trigger the clock policy to select pages from any process
- On a page fault
  - Take a page from the free list

### WSClock

- Wherever you need to reclaim a page
  - 1. Examine the PTE pointed to by clock hand.
  - 2. If reference bit is set
    - 1. Clear reference bit;
    - 2. Advance clock hand;
    - 3. Goto Done.
  - 3. If reference bit is not set
    - 1. If the timestamp of the PTE is older than a threshold
      - 1. Write the page to disk if it's dirty and use this page
      - 2. Goto Done
    - 2. Otherwise
      - 1. Advance clock hand
      - 2. Goto 1.
  - 4. Done
  - 5. If no victim page is chosen, randomly pick one

### The impact of WSClock

 One of the most important page replacement policies in practice



### Announcement

- No lecture this Thursday relocate those 80 minutes to anytime you like before 2/11 11:59:00pm.
- Reading quizzes due next Tuesday
- Start working on your project **ASAP**

Computer Science & Engineering





