



# last time (1)

kill() signal sending timing

user ID/group IDs

used by kernel to determine what to do

tracked for every process

libraries/utilities map to names

permission checks in system call handlers

chmod permissions

user ID (owner) / one group ID / others — read/write/exec

access control list

list of users/groups — read/write/exec for each

## last time (2)

user ID 0 ('root', 'superuser') — passes all permission checks

login program: runs as user ID 0

can access password database because user ID 0

set-user-ID programs

special bit says “run program with owner’s user ID”

system administrator can setup program to only do ‘safe’ things

example: sudo: allow only users config file to do things as root

example: allow users to shutdown only if no one logged in

system tracks addt'l user ID to help with those checks

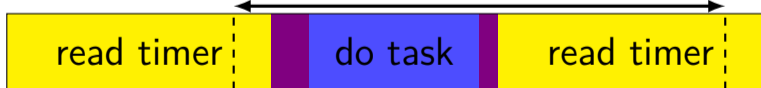
## anonymous feedback (1)

“Could we have some more office hours at the beginning/middle of the week instead of the end? I feel that it would be more helpful in the case that we have questions concerning the homework. Thank you.”

## anonymous feedback (2)

“Hello, I feel that the following would be helpful for the entire class as they start/continue to work on the homework: Is it possible for you to explain why the overhead time with `clock_gettime()` is longer than some of the provided scenarios? Even after incorporating the tips of Section 1.3, it still is longer. Thank you for your help.”

could be: system slow during overhead measurement, fast otherwise  
also can be task (e.g. empty function call) being optimized away  
also I think some are confused about what overhead is:



want to time task part, but timed measured includes extra stuff

## anonymous feedback (3)

“Do you mind going over what exactly the read, write, and execute permissions allow?”

regular files:

read — open file for reading

write — open file to write/print (modify contents)

execute — run the file as a program

directories:

read — list directory contents write — add/move/rename files in

directory execute (search) — access file/subdirectory within directory (if already know name)

## quiz Q2

question: when will control-C terminate program?

answer: when signal handler not setup yet

once `sigaction()` called, signal handler remains setup until another `sigaction/`etc. call

## quiz Q4

aaa1a/ccc1c: read-only

bbb1b: write3

make bbb1b owner, give owner rw permissions

have aaa1a+ccc1c (and maybe also bbb1b) in group

associate that group w/file + give group r permission only

make default permissions nothing



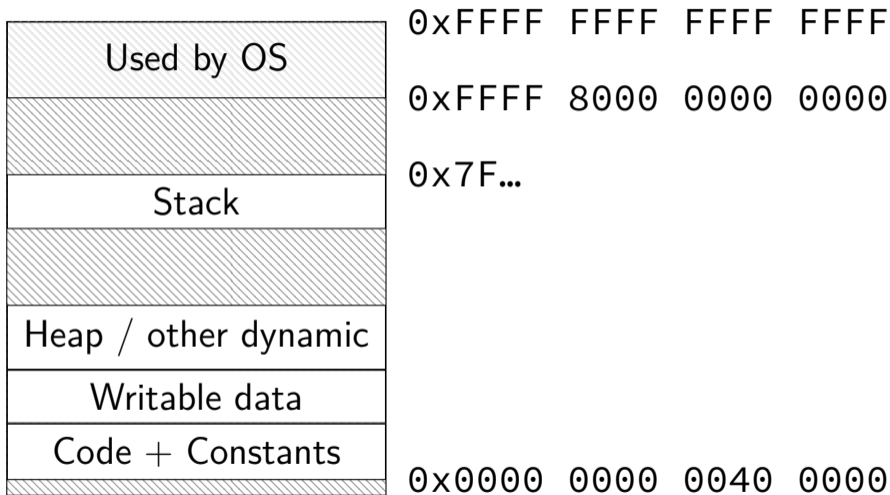
## anonymous feedback (4)

“Do you think you could post the recordings on Panopto please?  
The video player doesn't allow us to have captions or skip/go back  
10/sec intervals”

done

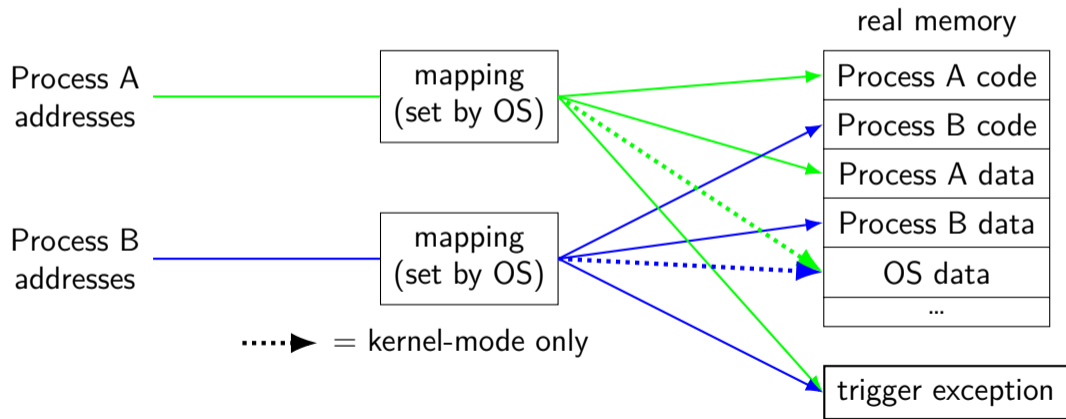
also should be importing panopto automatic [not great] captions into  
main viewer

# program memory



# address spaces

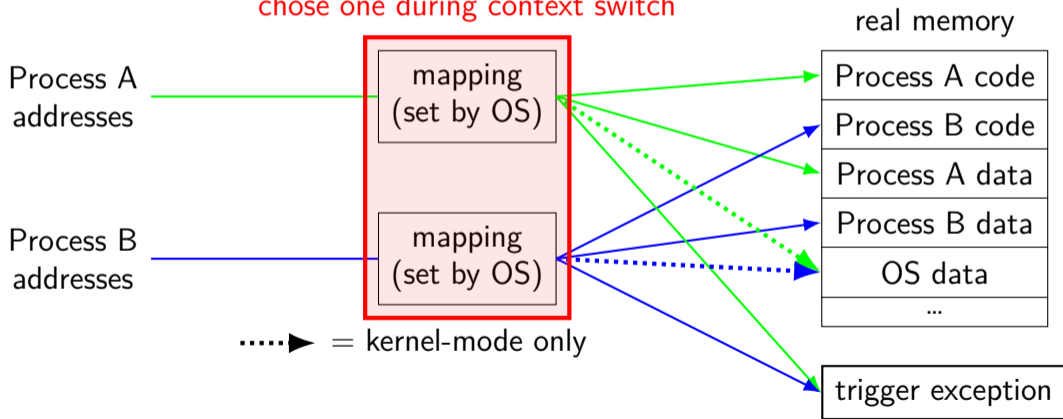
illusion of **dedicated memory**



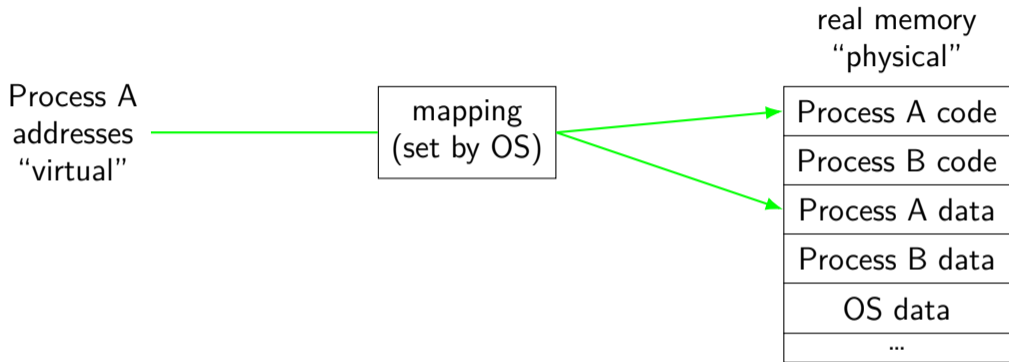
# address spaces

illusion of **dedicated memory**

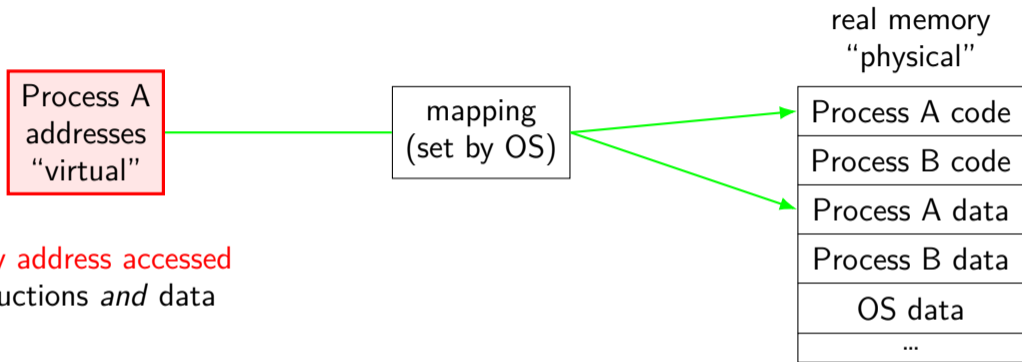
chose one during context switch



# address translation

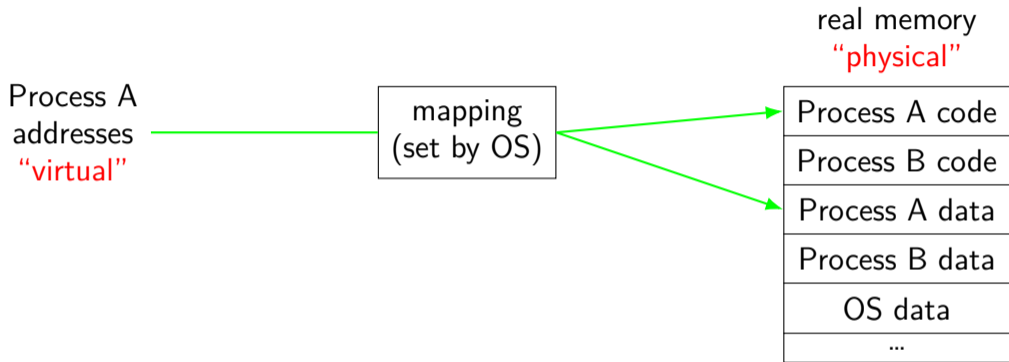


# address translation



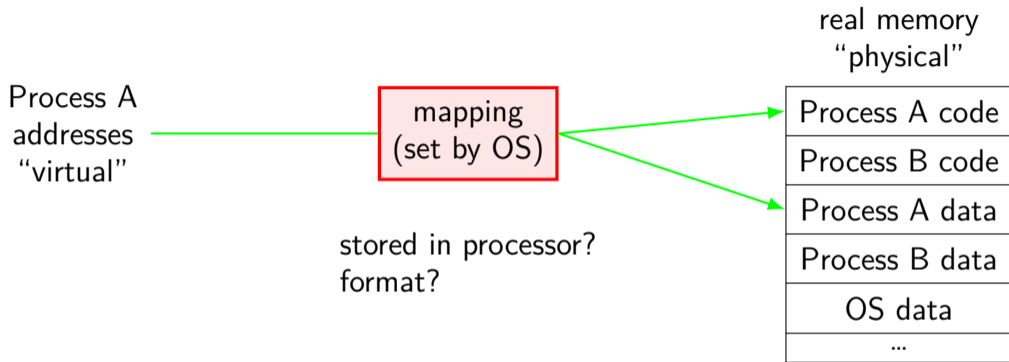
every address accessed  
instructions *and* data

# address translation



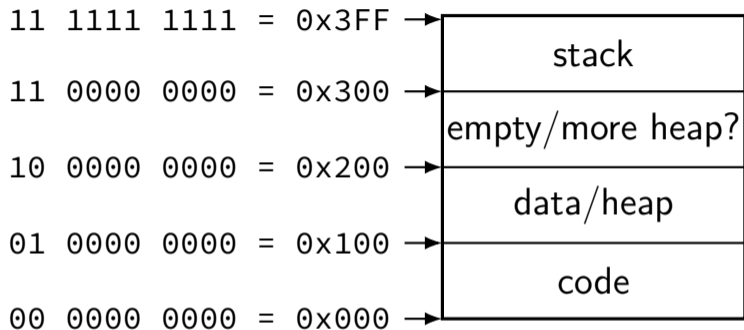
program addresses are 'virtual'  
real addresses are 'physical'  
can be **different sizes!**

# address translation

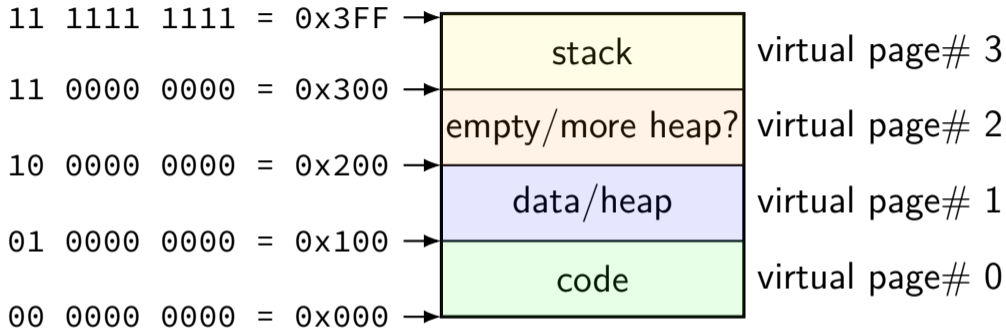




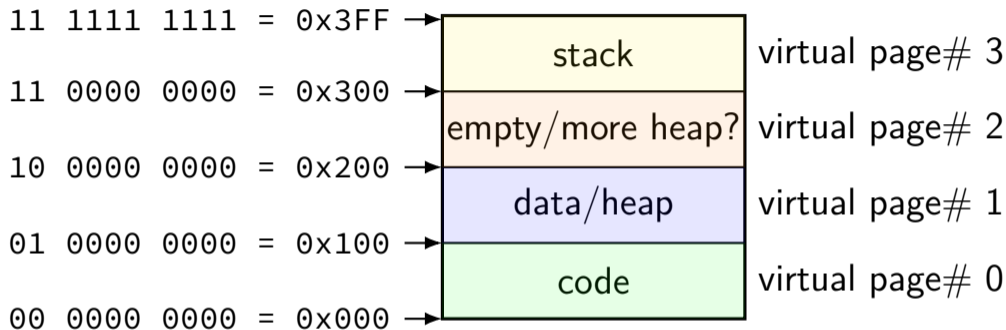
# toy program memory



# toy program memory

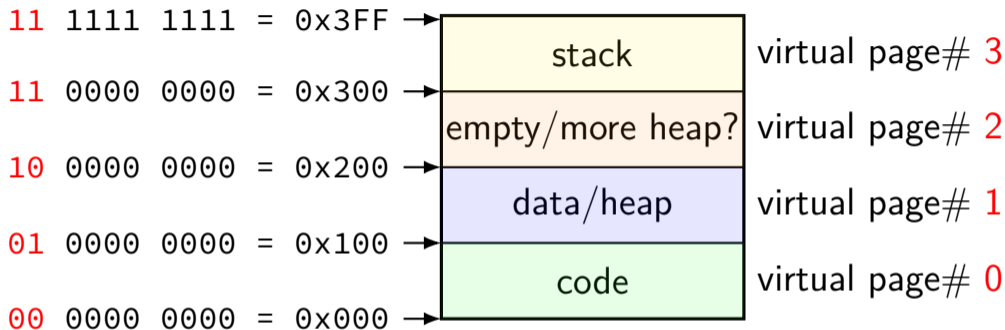


# toy program memory



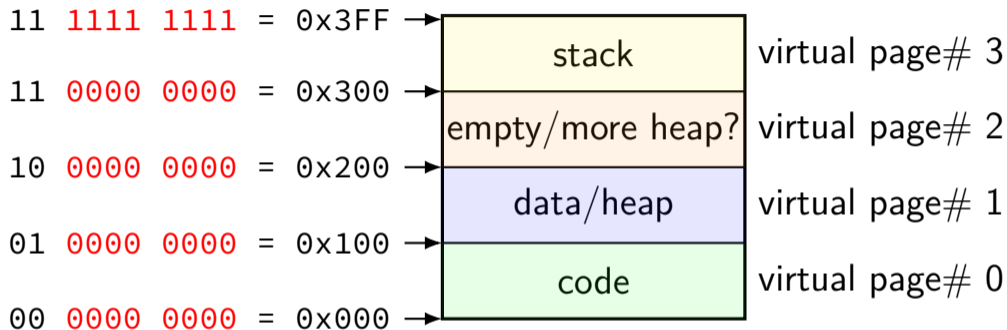
divide memory into **pages** ( $2^8$  bytes in this case)  
“virtual” = addresses the program sees

# toy program memory



page number is upper bits of address  
(because page size is power of two)

# toy program memory



rest of address is called **page offset**

# toy physical memory

program memory  
virtual addresses

11 0000 0000 to 11 1111 1111
10 0000 0000 to 10 1111 1111
01 0000 0000 to 01 1111 1111
00 0000 0000 to 00 1111 1111

real memory  
physical addresses

111 0000 0000 to 111 1111 1111
001 0000 0000 to 001 1111 1111
000 0000 0000 to 000 1111 1111

# toy physical memory

program memory  
virtual addresses

11 0000 0000 to 11 1111 1111
10 0000 0000 to 10 1111 1111
01 0000 0000 to 01 1111 1111
00 0000 0000 to 00 1111 1111

real memory  
physical addresses

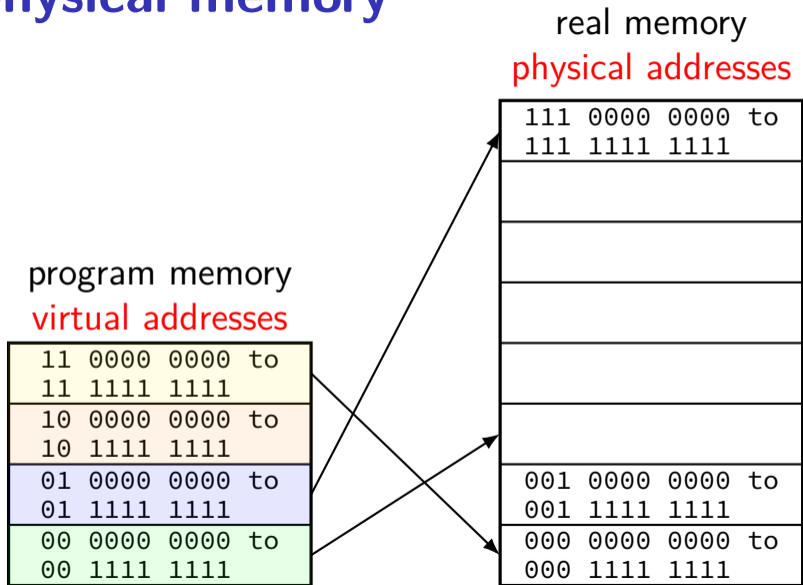
111 0000 0000 to 111 1111 1111
001 0000 0000 to 001 1111 1111
000 0000 0000 to 000 1111 1111

physical page 7

physical page 1

physical page 0

# toy physical memory





# toy physical memory

virtual page #	physical page #
00	010 (2)
01	111 (7)
10	<i>none</i>
11	000 (0)

program memory  
virtual addresses

11 0000 0000 to 11 1111 1111
10 0000 0000 to 10 1111 1111
01 0000 0000 to 01 1111 1111
00 0000 0000 to 00 1111 1111

real memory

physical addresses

111 0000 0000 to 111 1111 1111
001 0000 0000 to 001 1111 1111
000 0000 0000 to 000 1111 1111

# toy physical memory

virtual page #	physical page #
00	010 (2)
01	111 (7)
10	<i>none</i>
11	000 (0)

program memory  
virtual addresses

11 0000 0000 to
11 1111 1111
10 0000 0000 to
10 1111 1111
01 0000 0000 to
01 1111 1111
00 0000 0000 to
00 1111 1111

page  
table! real memory  
physical addresses

111 0000 0000 to
111 1111 1111
001 0000 0000 to
001 1111 1111
000 0000 0000 to
000 1111 1111

# toy page table lookup

virtual  
page #    valid?    physical page #

00	1	010 (2, code)
01	1	111 (7, data)
10	0	??? (ignored)
11	1	000 (0, stack)

# toy page table lookup

01 1101 0010 — address from CPU

virtual  
page #    valid?    physical page #

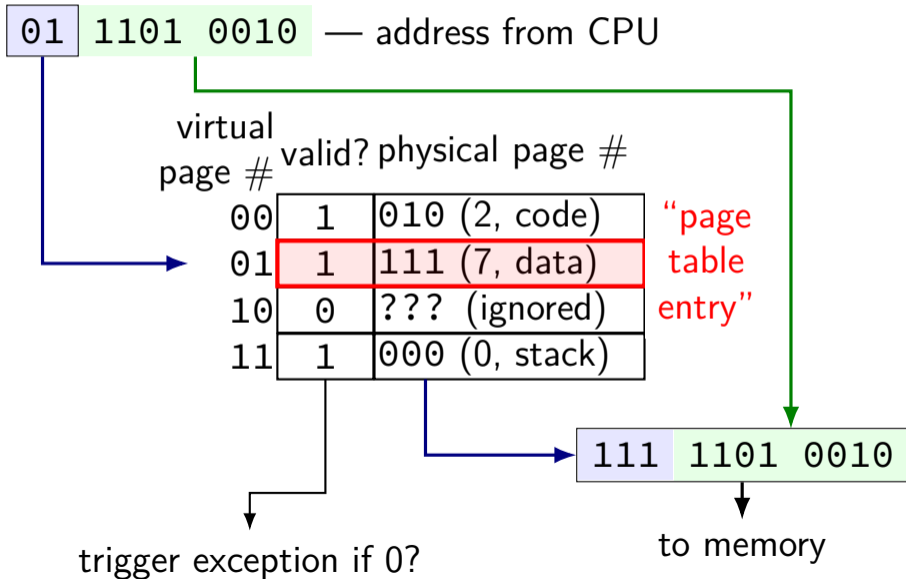
00	1	010 (2, code)
01	1	111 (7, data)
10	0	??? (ignored)
11	1	000 (0, stack)

111 1101 0010

trigger exception if 0?

to memory

# toy page table lookup



# t “virtual page number” lookup

01 1101 0010 — address from CPU

virtual  
page # valid? physical page #

00	1	010 (2, code)
01	1	111 (7, data)
10	0	??? (ignored)
11	1	000 (0, stack)

111 1101 0010

trigger exception if 0?

to memory

# toy page table lookup

01 1101 0010 — address from CPU

virtual  
page # valid? physical page #

00	1	010 (2, code)
01	1	111 (7, data)
10	0	??? (ignored)
11	1	000 (0, stack)

“physical page number”

111 1101 0010

trigger exception if 0?

to memory

# toy pag "page offset" lookup

01 1101 0010 — address from CPU

virtual  
page # valid? physical page #

00	1	010 (2, code)
01	1	111 (7, data)
10	0	??? (ignored)
11	1	000 (0, stack)

"page offset"

111 1101 0010

trigger exception if 0?

to memory



## on virtual address sizes

virtual address size = size of pointer?

often, but — sometimes part of pointer not used

example: typical x86-64 only use 48 bits

rest of bits have fixed value

virtual address size is amount used for mapping

# address space sizes

amount of stuff that can be addressed = address space size  
based on number of unique addresses

e.g. 32-bit virtual address =  $2^{32}$  byte virtual address space

e.g. 20-bit physical address =  $2^{20}$  byte physical address space

# address space sizes

amount of stuff that can be addressed = address space size  
based on number of unique addresses

e.g. 32-bit virtual address =  $2^{32}$  byte virtual address space

e.g. 20-bit physical address =  $2^{20}$  byte physical address space

what if my machine has 3GB of memory (not power of two)?

not all addresses in physical address space are useful

most common situation (since CPUs support having a lot of memory)

## exercise: page counting

suppose 32-bit virtual (program) addresses

and each page is 4096 bytes ( $2^{12}$  bytes)

how many virtual pages?

## exercise: page counting

suppose 32-bit virtual (program) addresses

and each page is 4096 bytes ( $2^{12}$  bytes)

how many virtual pages?

$$2^{32} / 2^{12} = 2^{20}$$

## exercise: page table size

suppose 32-bit virtual (program) addresses

suppose 30-bit physical (hardware) addresses

each page is 4096 bytes ( $2^{12}$  bytes)

page table entries have physical page #, valid bit, bit

how big is the page table (if laid out like ones we've seen)?

## exercise: page table size

suppose 32-bit virtual (program) addresses

suppose 30-bit physical (hardware) addresses

each page is 4096 bytes ( $2^{12}$  bytes)

page table entries have physical page #, valid bit, bit

how big is the page table (if laid out like ones we've seen)?

$2^{20}$  entries  $\times$  (18 + 1) bits per entry

issue: where can we store that?

## exercise: address splitting

and each page is 4096 bytes ( $2^{12}$  bytes)

split the address 0x12345678 into page number and page offset:



## exercise: address splitting

and each page is 4096 bytes ( $2^{12}$  bytes)

split the address 0x12345678 into page number and page offset:

page #: 0x12345; offset: 0x678

# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15)	physical page # (bits 4–14)	other bits and/or unused (bit 0-3)
----------------	-----------------------------	------------------------------------

# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15)	physical page # (bits 4–14)	other bits and/or unused (bit 0-3)
----------------	-----------------------------	------------------------------------

page table  
base register

0x00010000
------------

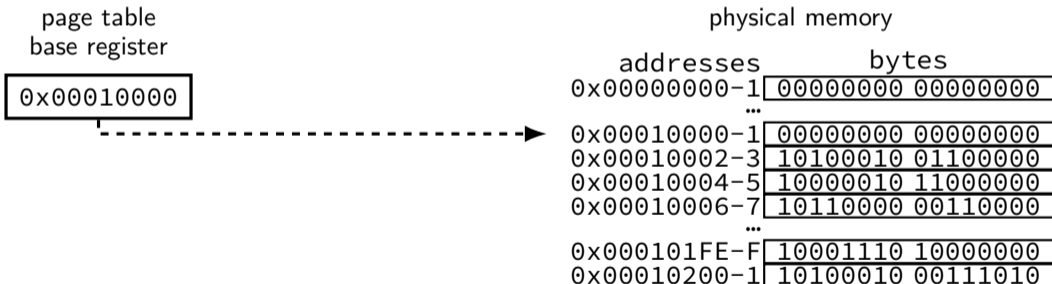


# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15)	physical page # (bits 4–14)	other bits and/or unused (bit 0-3)
----------------	-----------------------------	------------------------------------

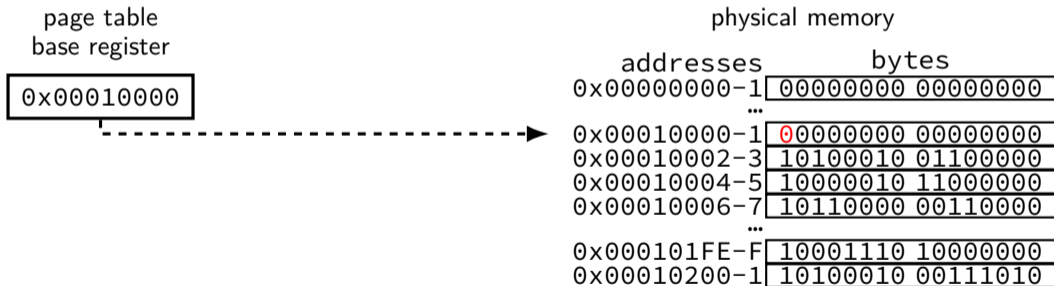


# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15)	physical page # (bits 4–14)	other bits and/or unused (bit 0-3)
----------------	-----------------------------	------------------------------------

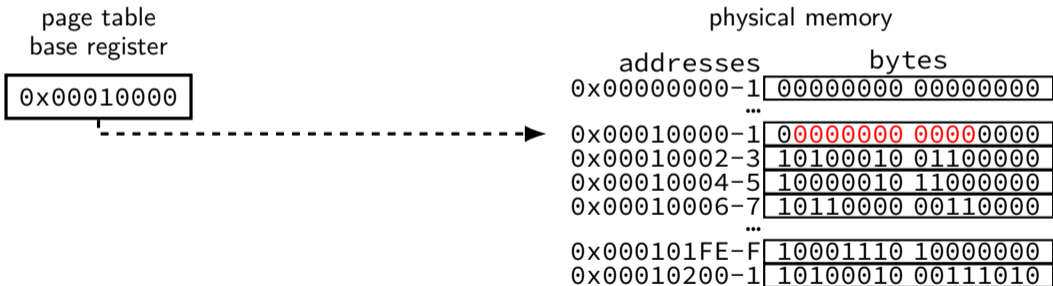


# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15)	physical page # (bits 4-14)	other bits and/or unused (bit 0-3)
----------------	-----------------------------	------------------------------------

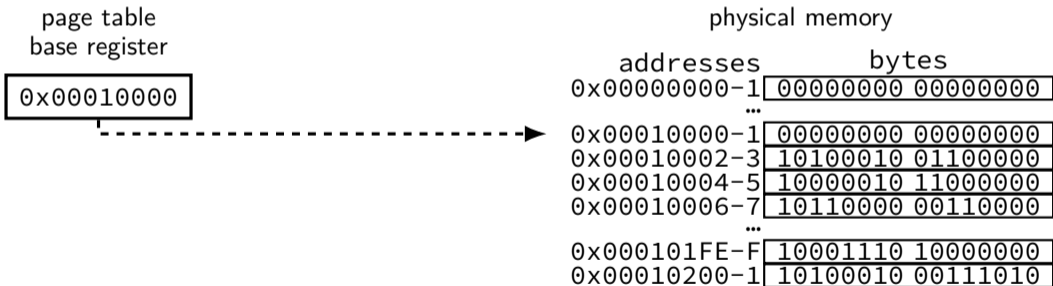


# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15)	physical page # (bits 4-14)	other bits and/or unused (bit 0-3)
----------------	-----------------------------	------------------------------------



# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15)	physical page # (bits 4–14)	other bits and/or unused (bit 0-3)
----------------	-----------------------------	------------------------------------

page table  
base register

0x00010000
------------

page table (logically)

virtual page #	valid?	...	physical page #
0000 0000	0	...	00 0000 0000
0000 0001	1	...	10 0010 0110
0000 0010	1	...	00 0000 1100
0000 0011	1	...	11 0000 0011
...			
1111 1111	1	...	00 1110 1000

physical memory

addresses	bytes
0x00000000-1	00000000 00000000
...	
0x00010000-1	00000000 00000000
0x00010002-3	10100010 01100000
0x00010004-5	10000010 11000000
0x00010006-7	10110000 00110000
...	
0x000101FE-F	10001110 10000000
0x00010200-1	10100010 00111010

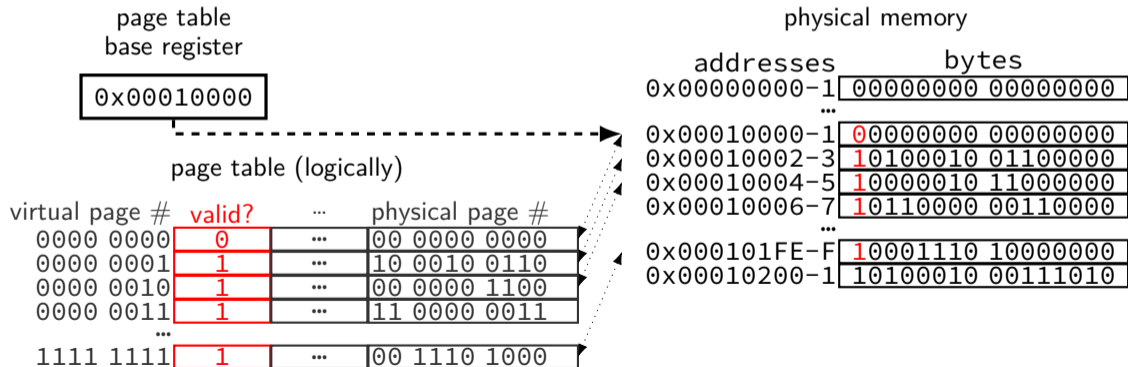


# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

**valid (bit 15)** | physical page # (bits 4–14) | other bits and/or unused (bit 0-3)

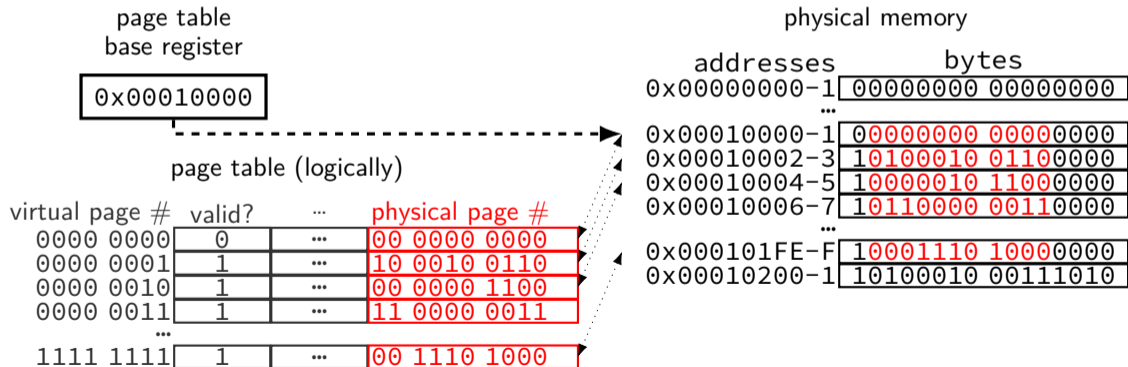


# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15) **physical page # (bits 4–14)** other bits and/or unused (bit 0-3)

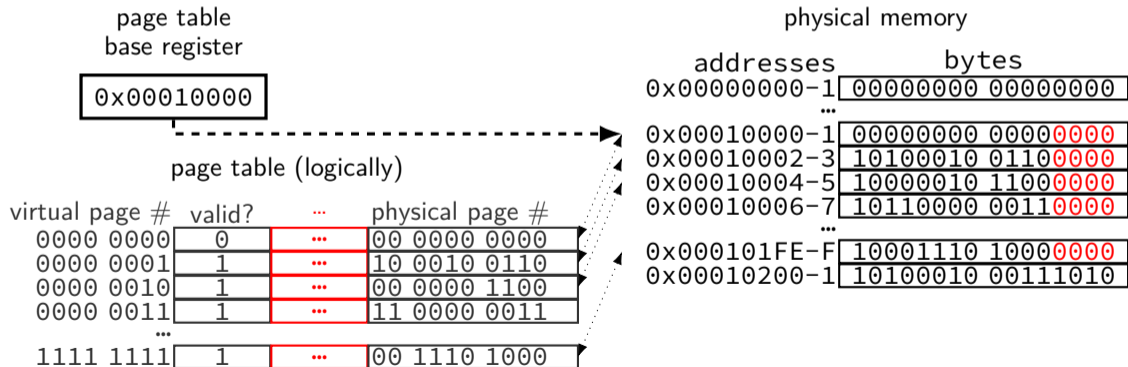


# page tables in memory

where can processor store megabytes of page tables? **in memory**

page table entry layout (chosen by processor)

valid (bit 15) | physical page # (bits 4–14) | **other bits and/or unused (bit 0-3)**

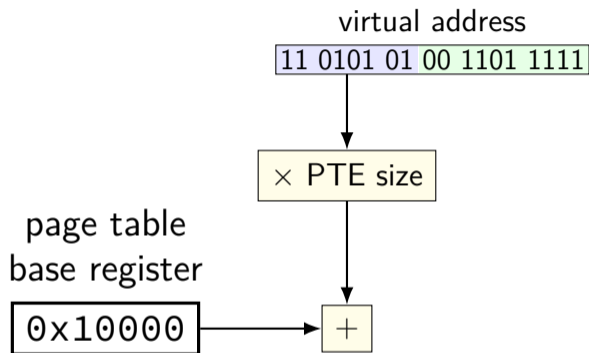


# memory access with page table

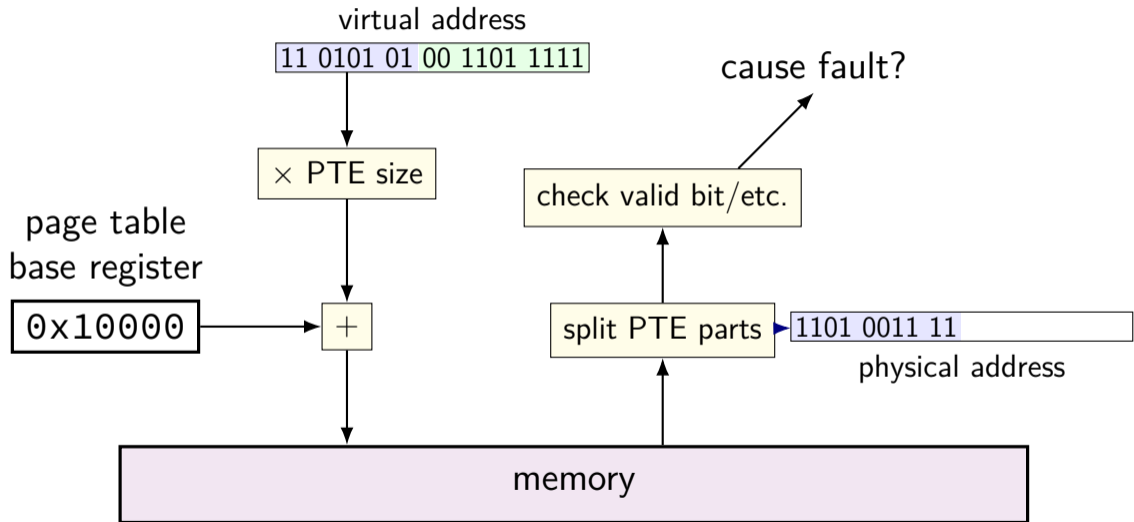
virtual address

11 0101 01 00 1101 1111

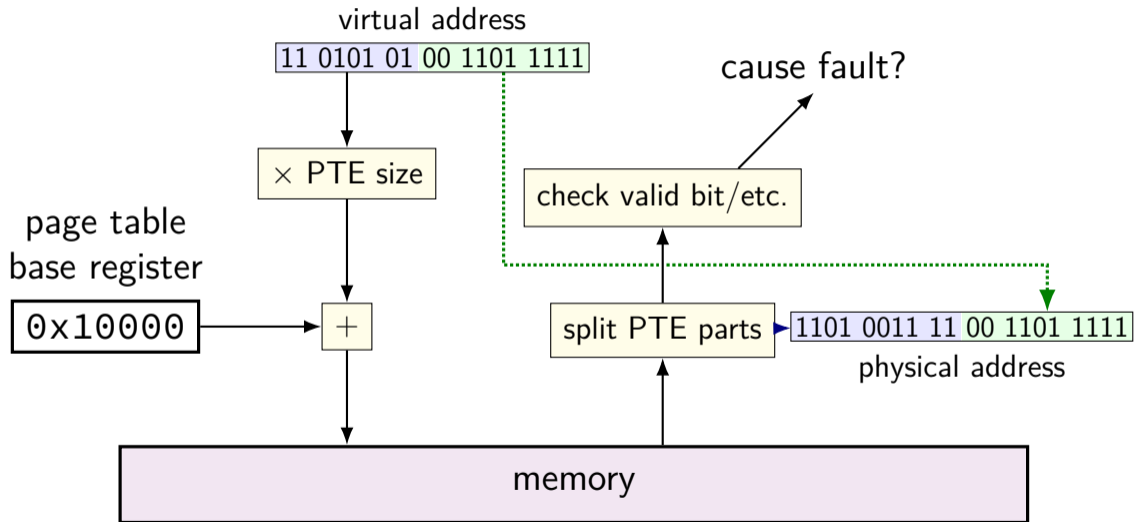
# memory access with page table



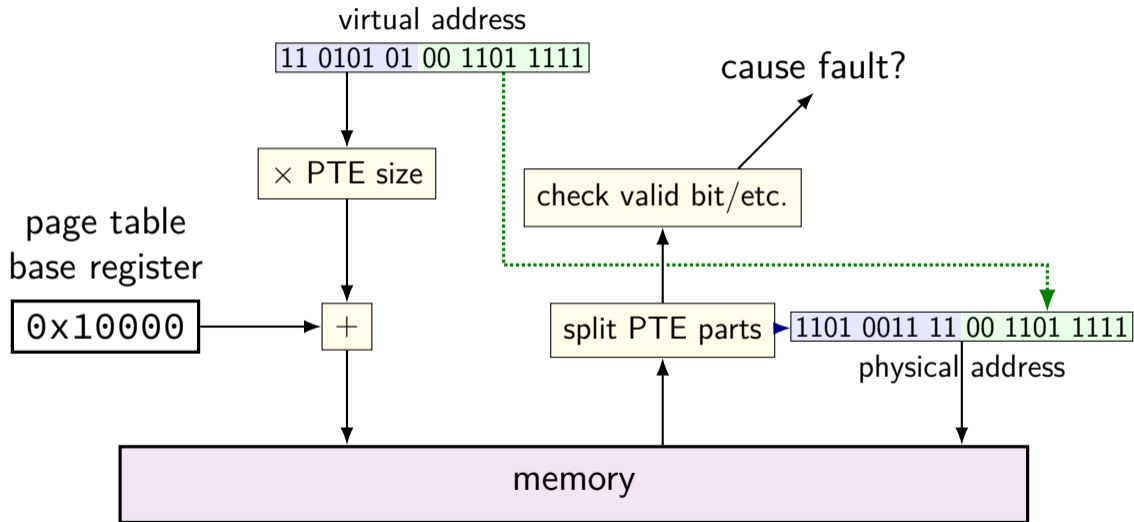
# memory access with page table



# memory access with page table

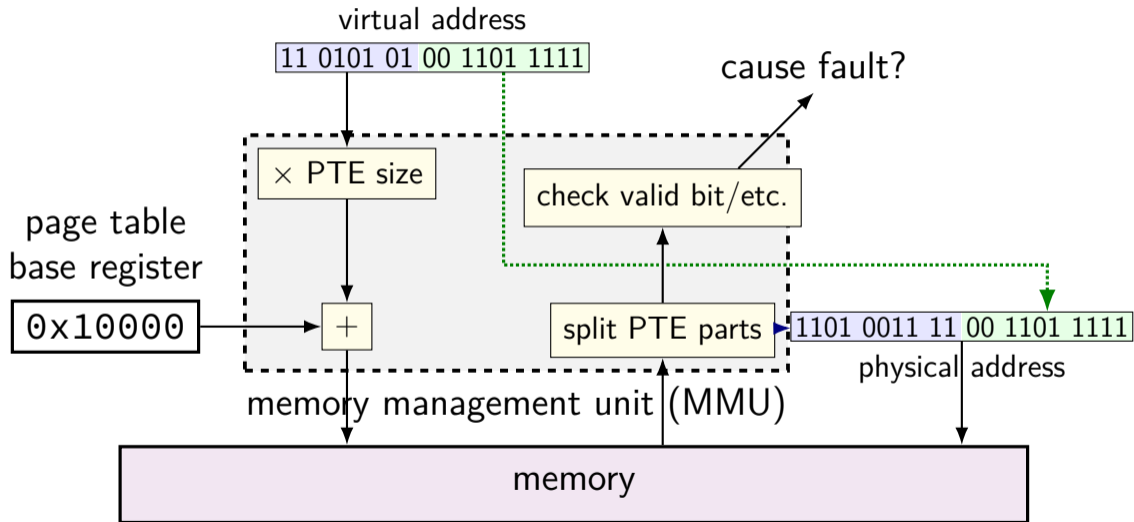


# memory access with page table

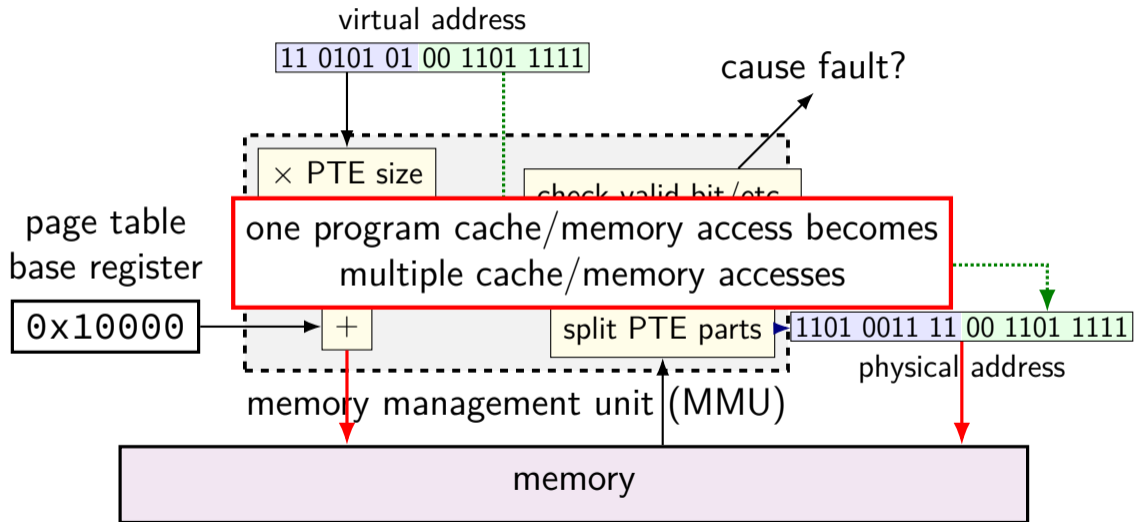




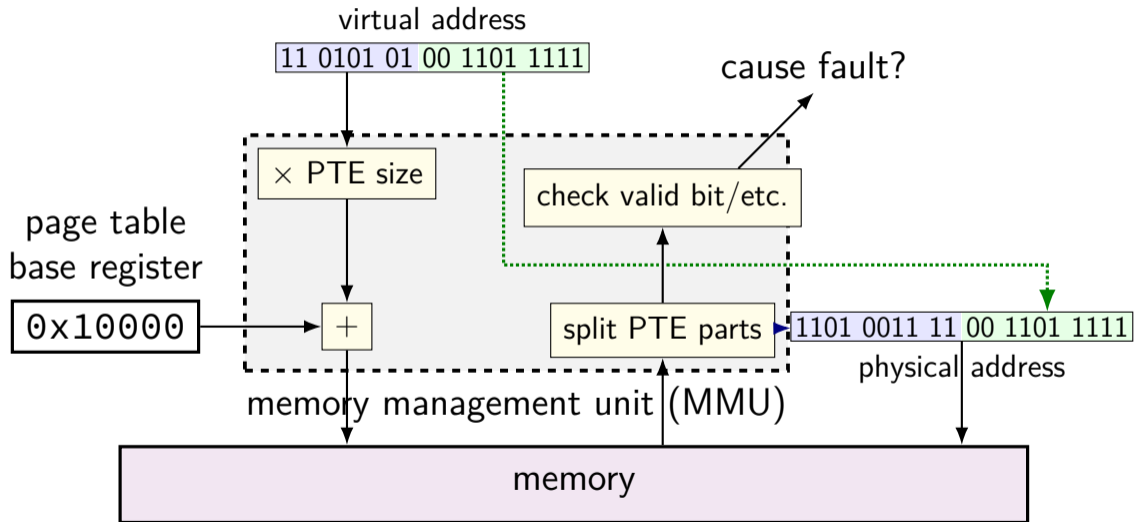
# memory access with page table



# memory access with page table



# memory access with page table



# backup slides