


Virtual Memory: Part 1


Computer Architecture

Instructors:

Charles Reiss and Daniel Graham

address translation



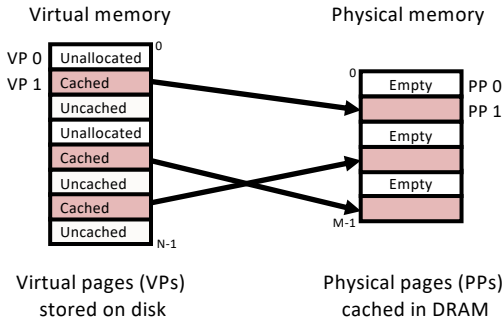


VM as a Tool for Caching

Conceptually, *virtual memory* is an array of N contiguous bytes stored on disk.

The contents of the array on disk are cached in *physical memory (DRAM cache)*

These cache blocks are called *pages*




DRAM Cache Organization

DRAM cache organization driven by the enormous miss penalty

DRAM is about **10x** slower than SRAM

Disk is about **10,000x** slower than



DRAM Cache Organization

DRAM cache organization driven by the enormous miss penalty

DRAM is about **10x** slower than SRAM

Disk is about **10,000x** slower than DRAM

Consequences

Large page (block) size: typically 4 KB, sometimes 4 MB


Fully associative

Any Virtual Page can be placed in any Physical Page
Requires a “large” mapping function – different from memories

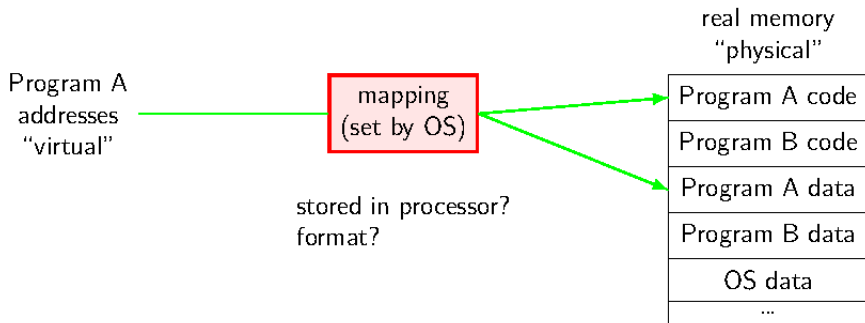
Highly sophisticated, expensive replacement algorithms

Too complicated and open-ended to be implemented in hardware

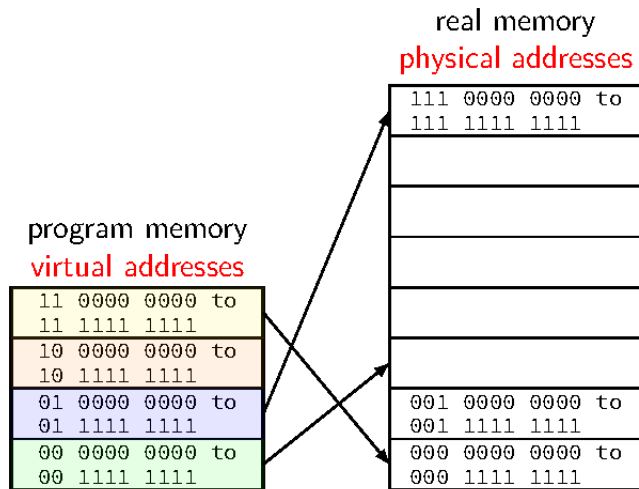
Write-back rather than write-through



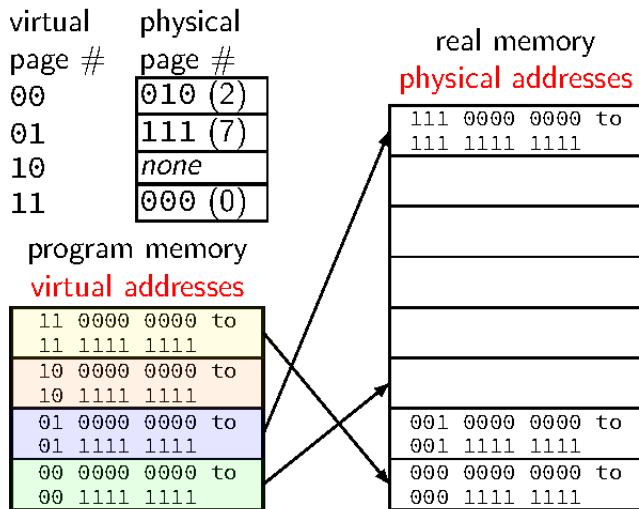
address translation



toy physical memory



toy physical memory



toy physical memory

page table!

virtual page #	physical page #
00	010 (2)
01	111 (7)
10	none
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

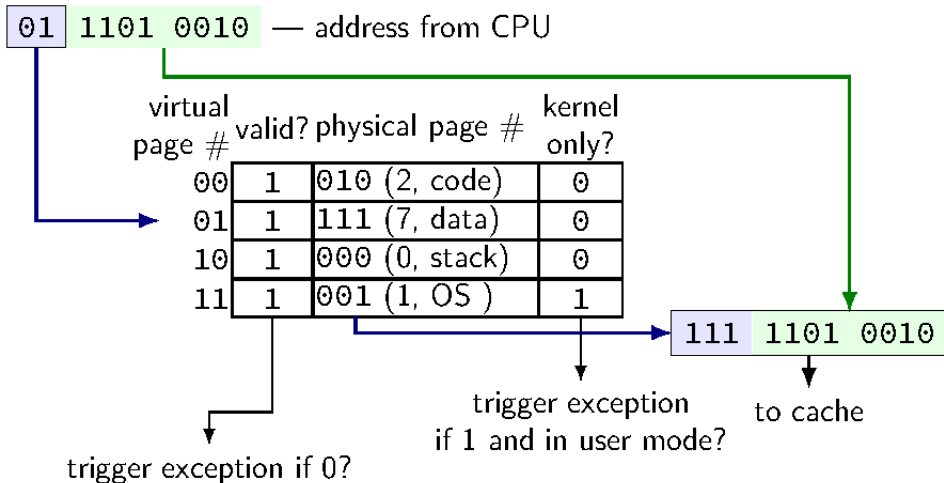
real memory
physical addresses

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

Address Translation

virtual page #	valid?	physical page #	kernel only?
00	1	010 (2, code)	0
01	1	111 (7, data)	0
10	1	000 (0, stack)	0
11	1	001 (1, OS)	1

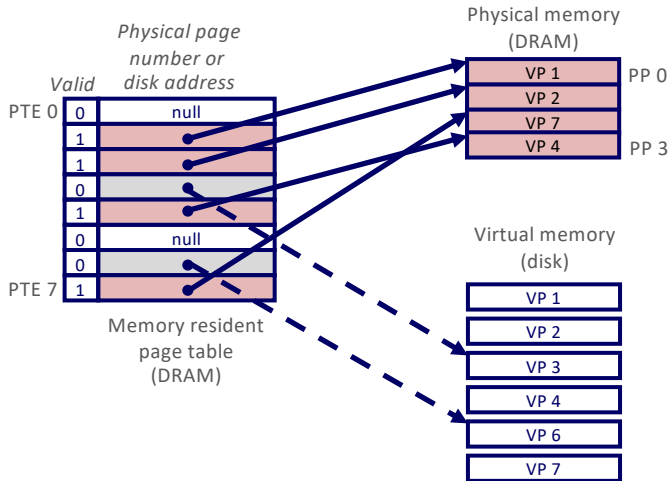
Address Translation



Enabling Data Structure: Page Table

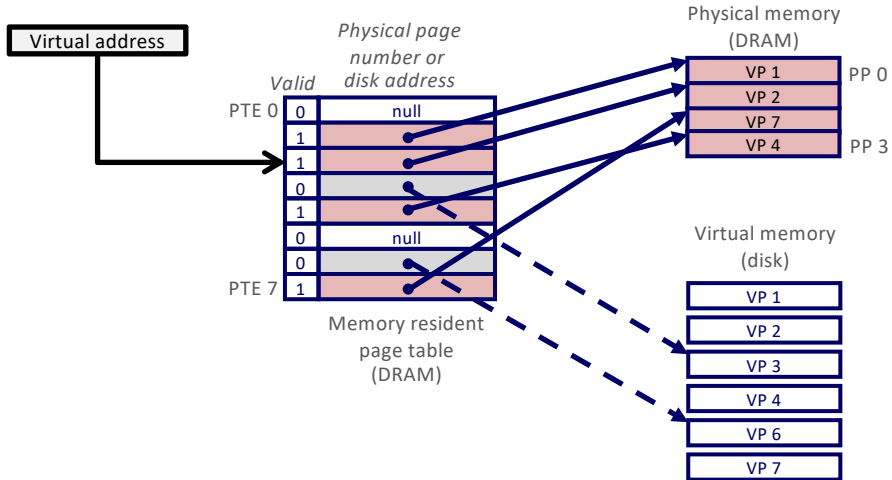
A *page table* is an array of page table entries (PTEs) that maps virtual pages to physical pages.

Per-process kernel data structure in DRAM



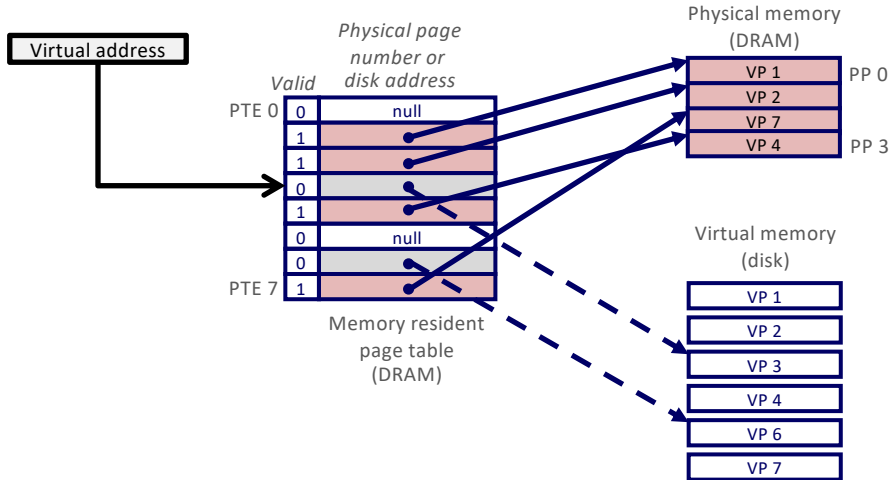
Page Hit

Page hit: reference to VM word that is in physical memory (DRAM cache hit)



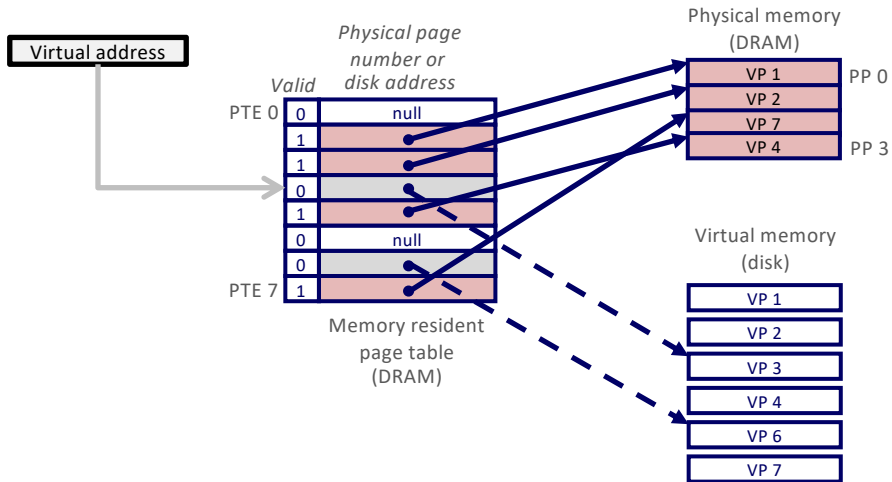
Page Fault

Page fault: reference to VM word that is not in physical memory (DRAM cache miss)



Handling Page Fault

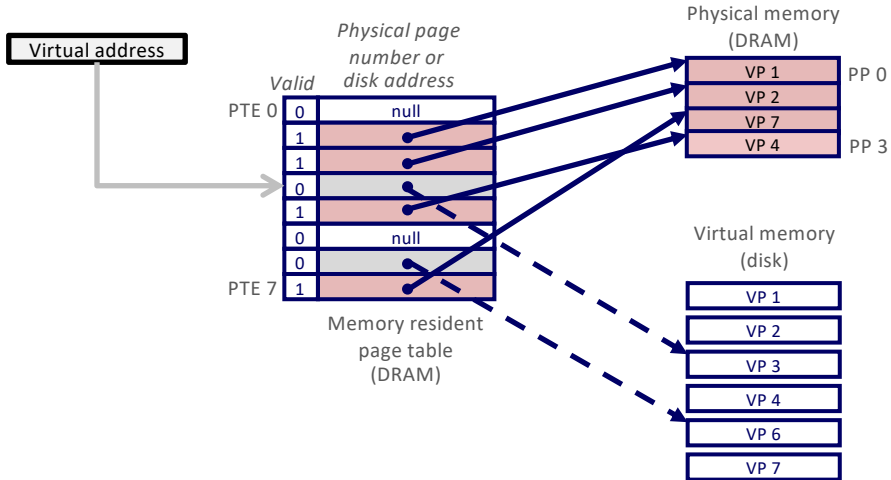
Page miss causes page fault (an exception)



Handling Page Fault

Page miss causes page fault (an exception)

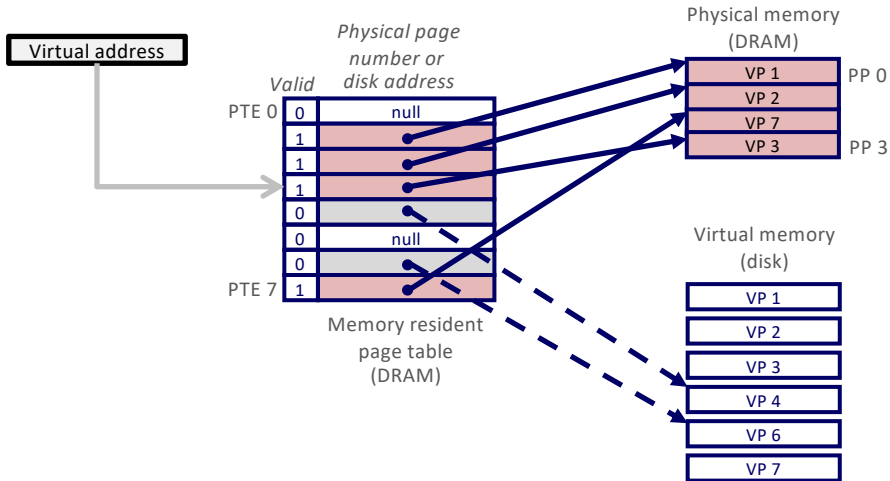
Page fault handler selects a victim to be evicted (here VP 4)



Handling Page Fault

Page miss causes page fault (an exception)

Page fault handler selects a victim to be evicted (here VP 4)

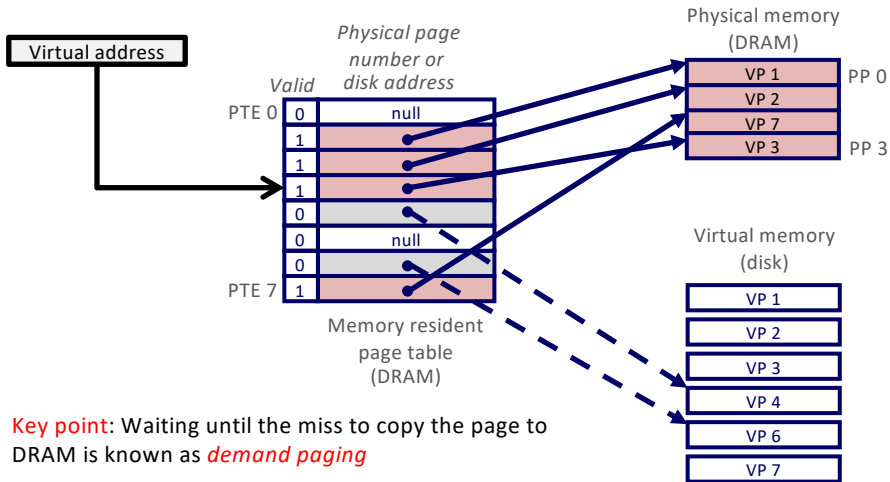


Handling Page Fault

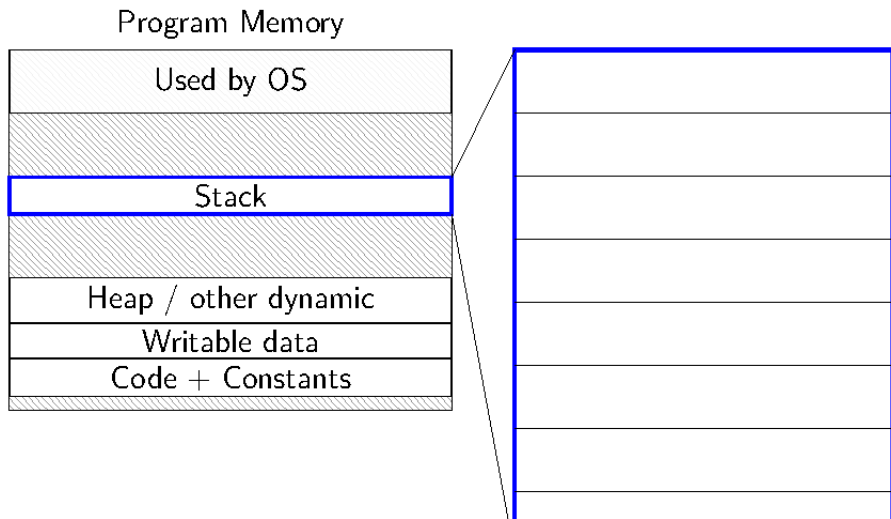
Page miss causes page fault (an exception)

Page fault handler selects a victim to be evicted (here VP 4)

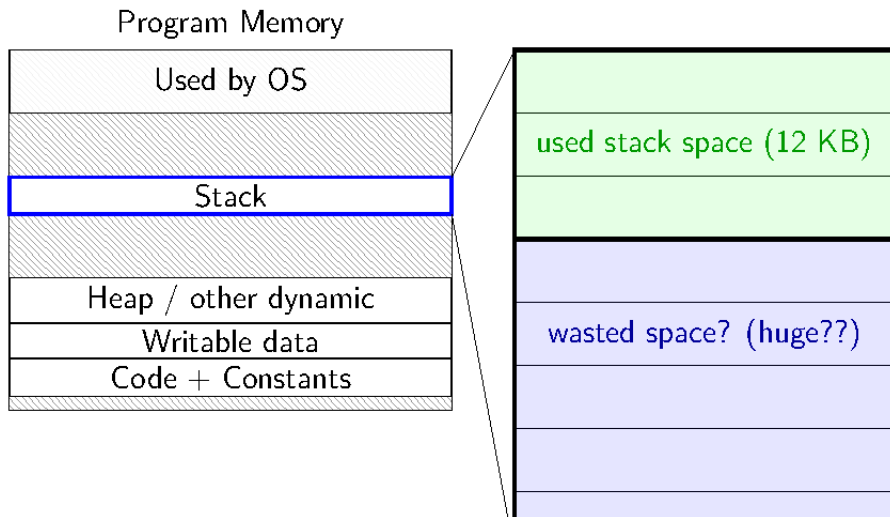
Offending instruction is restarted: page hit!



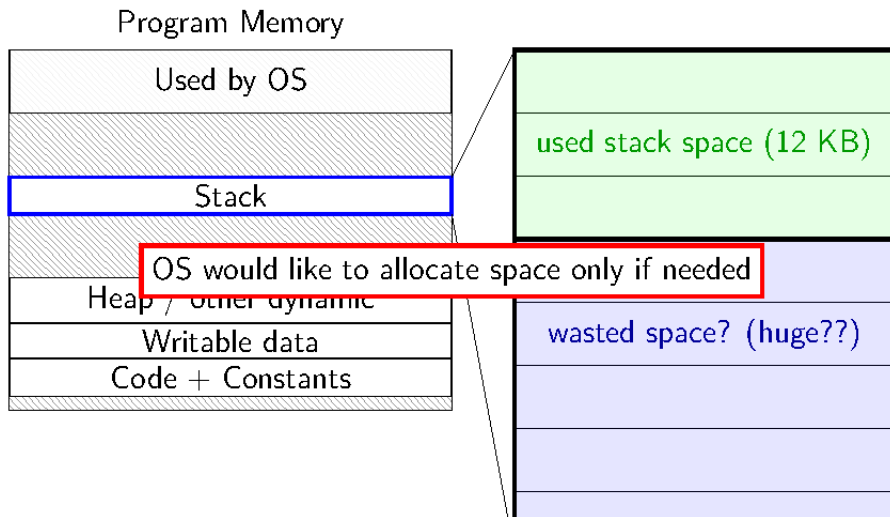
space on demand



space on demand



space on demand



allocating space on demand

`%rsp = 0x7FFFC000`

```
...  
// requires more stack space  
A: pushq %rbx  
  
B: movq 8(%rcx), %rbx  
C: addq %rbx, %rax  
...
```

VPN

```
...  
0x7FFFB  
0x7FFFC  
0x7FFFD  
0x7FFFE  
0x7FFFF  
...
```

valid? physical
page

valid?	physical page
...	...
0	---
1	0x200DF
1	0x12340
1	0x12347
1	0x12345
...	...

allocating space on demand

`%rsp = 0x7FFFC000`

```
...  
// requires more stack space  
A: pushq %rbx → page fault!  
B: movq 8(%rcx), %rbx  
C: addq %rbx, %rax  
...
```

VPN

```
...  
0x7FFFB  
0x7FFFC  
0x7FFFD  
0x7FFFE  
0x7FFFF  
...
```

valid? physical
page

valid?	physical page
...	...
0	---
1	0x200DF
1	0x12340
1	0x12347
1	0x12345
...	...

pushq triggers exception
hardware says "accessing address 0x7FFBFF8"
OS looks up what's should be there — "stack"

allocating space on demand

`%rsp = 0x7FFFC000`

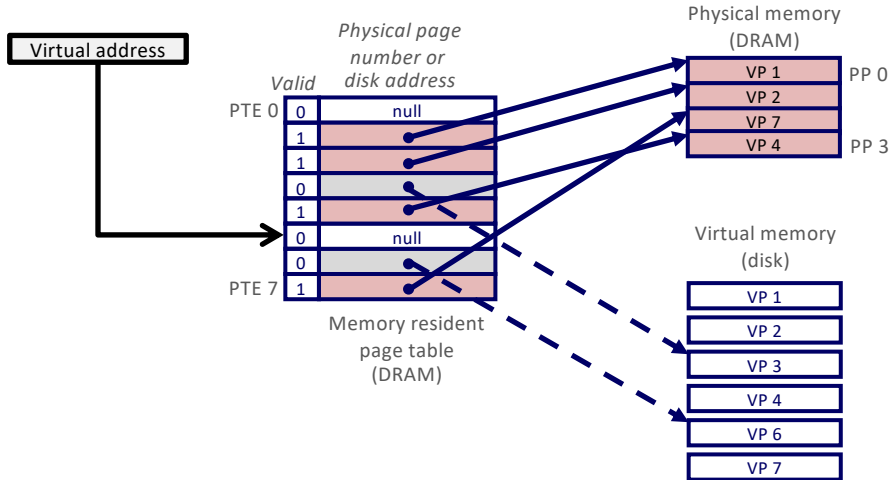
```
...  
// requires more stack space  
A: pushq %rbx restarted  
B: movq 8(%rcx), %rbx  
C: addq %rbx, %rax  
...
```

VPN	valid?	physical page
...
0x7FFFB	1	0x200D8
0x7FFFC	1	0x200DF
0x7FFFD	1	0x12340
0x7FFFE	1	0x12347
0x7FFFF	1	0x12345
...

in exception handler, OS allocates more stack space
OS updates the page table
then returns to retry the instruction

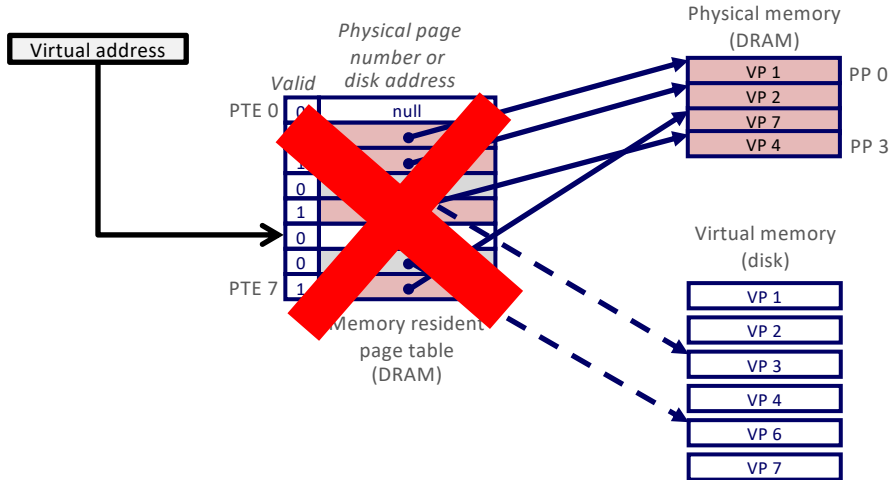
Page Fault

Page fault: reference to VM word that is not in physical memory (DRAM cache miss)



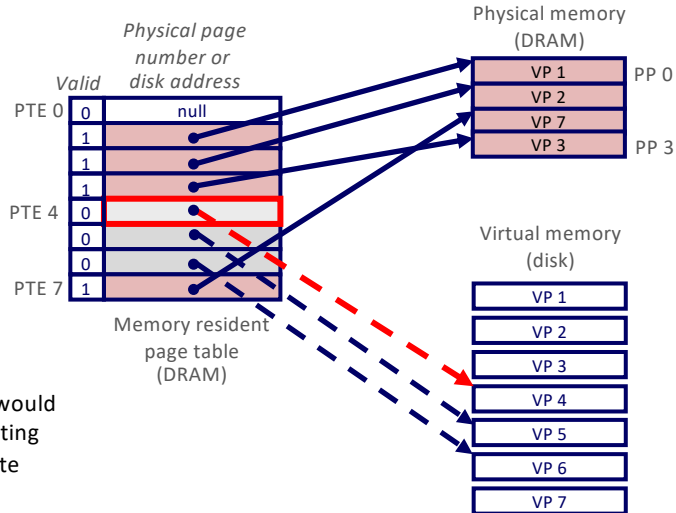
Page Fault

Page fault: reference to VM word that is not in physical memory (DRAM cache miss)



Allocating Pages

Allocating a new page (VP) of virtual memory.



A call to malloc would cause the Operating System to allocate Memory.

Locality to the Rescue Again!

Virtual memory seems terribly inefficient, but it works because of locality. *(On modern system often the program fits in DRAM because we have a lot of ram)*

At any point in time, programs tend to access a set of active virtual pages called the *working set*

Programs with better temporal locality will have smaller working sets

Locality to the Rescue Again!

If (working set size < main memory size)

Good performance for one process after compulsory misses

If (SUM(working set sizes) > main memory size)

Thrashing: Performance meltdown where pages are swapped (copied) in and out continuously

page tables in memory

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


page table entry layout

valid (bit 15)	kernel (bit 14)	physical page # (bits 4-13)	unused (bit 0-3)
----------------	-----------------	-----------------------------	------------------

page tables in memory

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

page table entry layout



page table base register


0x00010000




page tables in memory

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


page table entry layout



page table base register




physical memory



page tables in memory

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

page table entry layout



page table base register

0x00010000


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



page table base register

0x00010000

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



page table base register

0x00010000

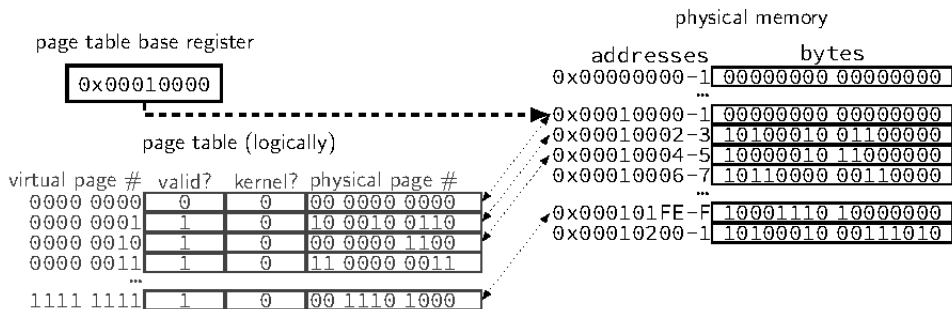

physical memory

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**

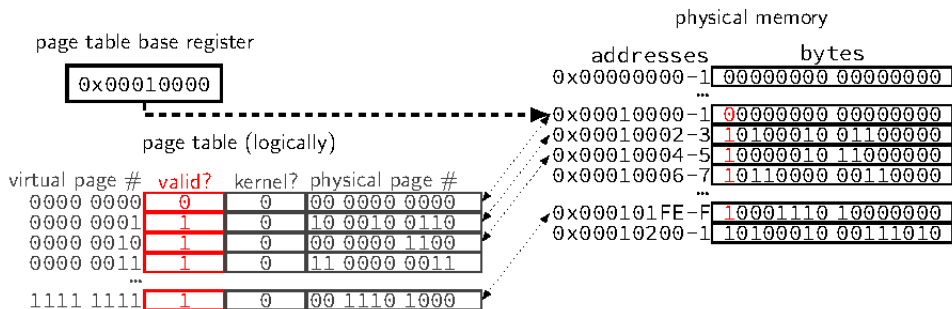

page table entry layout



page tables in memory

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

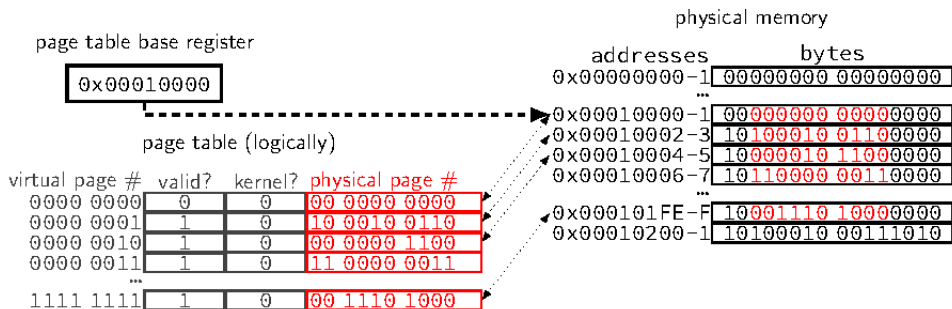
page table entry layout



page tables in memory

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



page table entry layout



page tables in memory

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

page table entry layout




memory access with page table

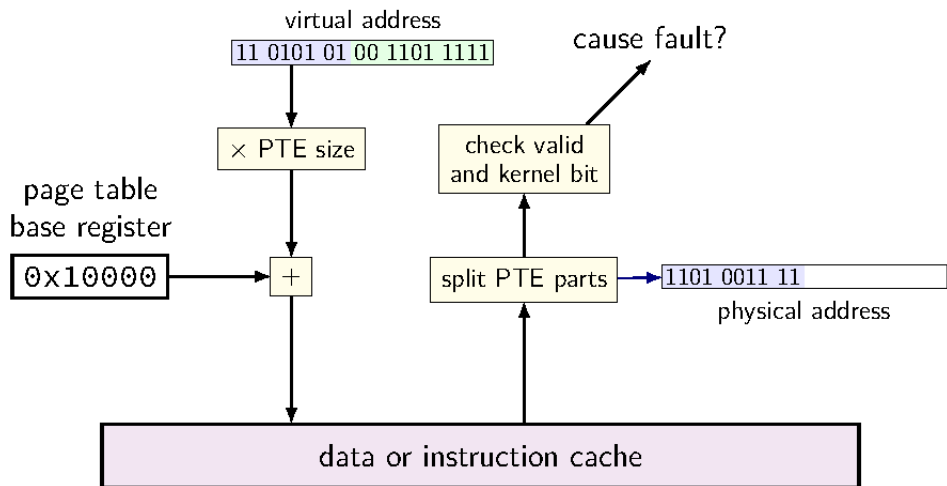
virtual address

11 0101 01 00 1101 1111

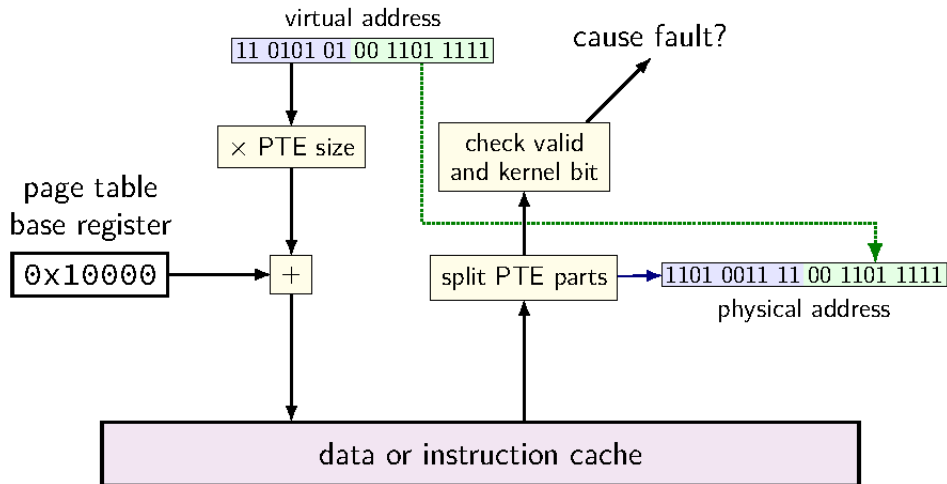
memory access with page table



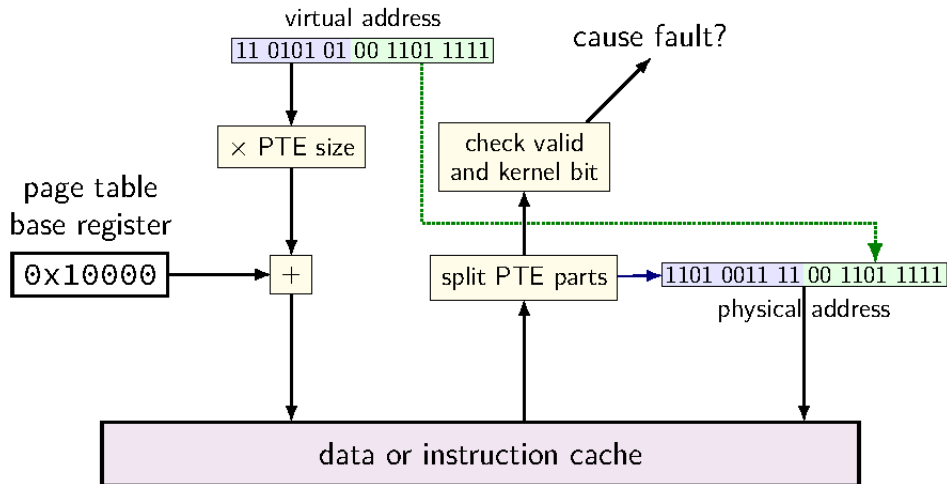
memory access with page table




memory access with page table



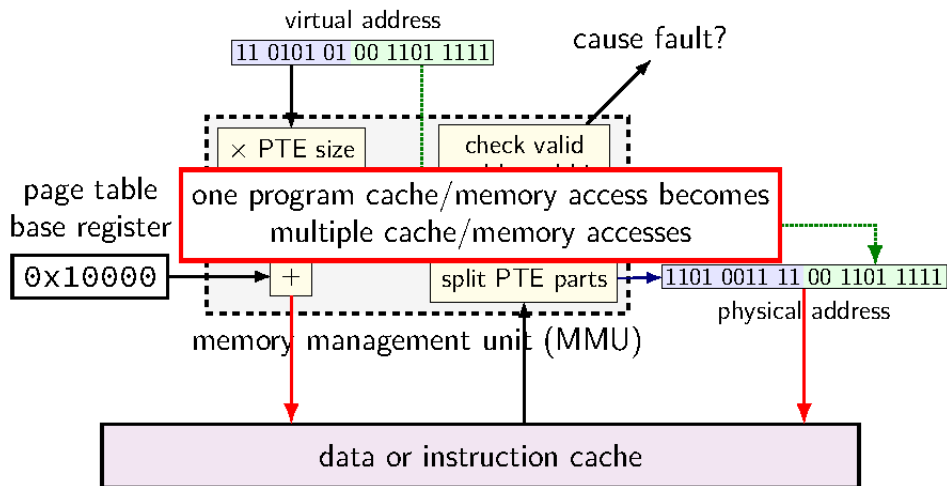
memory access with page table



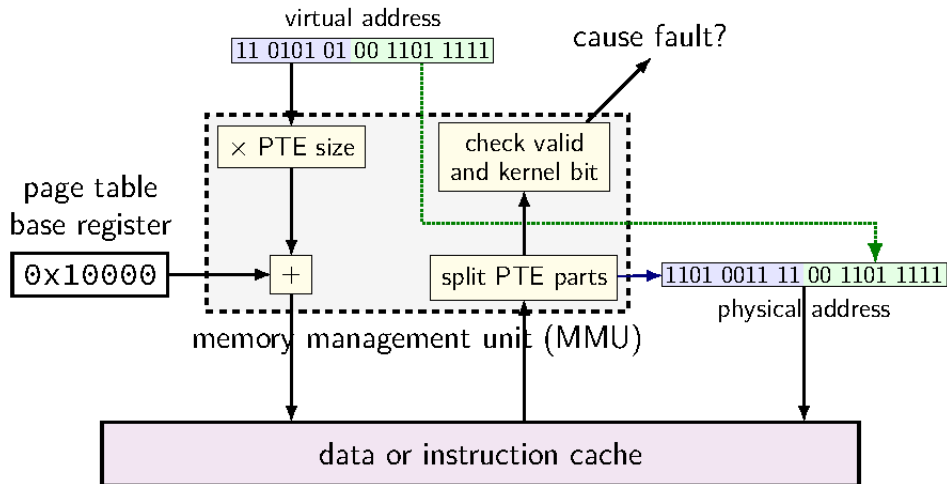
memory access with page table




memory access with page table



memory access with page table




MMUs in the pipeline



up to four memory accesses per instruction

MMUs in the pipeline



up to four memory accesses per instruction

challenging to make this fast (topic for a future date)

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

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, kernel-mode 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, kernel-mode bit

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

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

issue: where can we store that?

20 Mega Bits

\sim 2.5 Mega Bytes

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

two big problems to solve later

two **extra cache accesses** seems really slow

solution: more caches (called “TLB”, topic for later weeks)

page tables seem **really huge**

solution: not just a flat table

exercise setup

5-bit virtual addresses, 6-bit physical addresses, 8-byte pages

page table

virtual page #	valid?	physical page #
00	1	010
01	1	111
10	0	000
11	1	000

physical addresses	bytes
0x00-3	00 11 22 33
0x04-7	44 55 66 77
0x08-B	88 99 AA BB
0x0C-F	CC DD EE FF
0x10-3	1A 2A 3A 4A
0x14-7	1B 2B 3B 4B
0x18-B	1C 2C 3C 4C
0x1C-F	1C 2C 3C 4C

physical addresses	bytes
0x20-3	D0 D1 D2 D3
0x24-7	D4 D5 D6 D7
0x28-B	89 9A AB BC
0x2C-F	CD DE EF F0
0x30-3	BA 0A BA 0A
0x34-7	CB 0B CB 0B
0x38-B	DC 0C DC 0C
0x3C-F	EC 0C EC 0C

exercise setup

5-bit virtual addresses, 6-bit physical addresses, 8-byte pages

page table

virtual page #	valid?	physical page #
00	1	010
01	1	111
10	0	000
11	1	000

physical addresses	bytes
0x00-3	00 11 22 33
0x04-7	44 55 66 77
0x08-B	88 99 AA BB
0x0C-F	CC DD EE FF
0x10-3	1A 2A 3A 4A
0x14-7	1B 2B 3B 4B
0x18-B	1C 2C 3C 4C
0x1C-F	1C 2C 3C 4C

physical addresses	bytes
phys. page 0	1 D2 D3
	5 D6 D7
phys. page 1	A AB BC
	E EF F0
0x30-3	BA 0A BA 0A
0x34-7	CB 0B CB 0B
0x38-B	DC 0C DC 0C
0x3C-F	EC 0C EC 0C

exercise

5-bit virtual addresses, 6-bit physical addresses, 8-byte pages

(VAs) 0x18 = ???; 0x03 = ???; 0x0A = ???; 0x13 = ???

page table

virtual page #	valid?	physical page #
00	1	010
01	1	111
10	0	000
11	1	000

physical addresses	bytes
0x00-3	00 11 22 33
0x04-7	44 55 66 77
0x08-B	88 99 AA BB
0x0C-F	CC DD EE FF
0x10-3	1A 2A 3A 4A
0x14-7	1B 2B 3B 4B
0x18-B	1C 2C 3C 4C
0x1C-F	1C 2C 3C 4C

physical addresses	bytes
0x20-3	D0 D1 D2 D3
0x24-7	D4 D5 D6 D7
0x28-B	89 9A AB BC
0x2C-F	CD DE EF F0
0x30-3	BA 0A BA 0A
0x34-7	CB 0B CB 0B
0x38-B	DC 0C DC 0C
0x3C-F	EC 0C EC 0C

exercise

5-bit virtual addresses, 6-bit physical addresses, 8-byte pages

(VAs) $0x18 = (11000 \rightarrow 000000) 00$; $0x03 = ???$; $0x0A = ???$; $0x13 = ???$

page table

virtual page #	valid?	physical page #
00	1	010
01	1	111
10	0	000
11	1	000

physical addresses	bytes
0x00-3	00 11 22 33
0x04-7	44 55 66 77
0x08-B	88 99 AA BB
0x0C-F	CC DD EE FF
0x10-3	1A 2A 3A 4A
0x14-7	1B 2B 3B 4B
0x18-B	1C 2C 3C 4C
0x1C-F	1C 2C 3C 4C

physical addresses	bytes
0x20-3	D0 D1 D2 D3
0x24-7	D4 D5 D6 D7
0x28-B	89 9A AB BC
0x2C-F	CD DE EF F0
0x30-3	BA 0A BA 0A
0x34-7	CB 0B CB 0B
0x38-B	DC 0C DC 0C
0x3C-F	EC 0C EC 0C

exercise

5-bit virtual addresses, 6-bit physical addresses, 8-byte pages

(VAs) $0x18 = (11000 \rightarrow 000000) 00$; $0x03 = (00011 \rightarrow 010011) 0x4A$;
 $0x0A = ???$; $0x13 = ???$

page table

virtual page #	valid?	physical page #
00	1	010
01	1	111
10	0	000
11	1	000

physical addresses	bytes
0x00-3	00 11 22 33
0x04-7	44 55 66 77
0x08-B	88 99 AA BB
0x0C-F	CC DD EE FF
0x10-3	1A 2A 3A 4A
0x14-7	1B 2B 3B 4B
0x18-B	1C 2C 3C 4C
0x1C-F	1C 2C 3C 4C

physical addresses	bytes
0x20-3	D0 D1 D2 D3
0x24-7	D4 D5 D6 D7
0x28-B	89 9A AB BC
0x2C-F	CD DE EF F0
0x30-3	BA 0A BA 0A
0x34-7	CB 0B CB 0B
0x38-B	DC 0C DC 0C
0x3C-F	EC 0C EC 0C

exercise

5-bit virtual addresses, 6-bit physical addresses, 8-byte pages

(VAs) $0x18 = (11000 \rightarrow 000000) 00$; $0x03 = (00011 \rightarrow 010011) 0x4A$;
 $0x0A = (01010 \rightarrow 111010) 0xDC$; $0x13 = ???$

page table

virtual page #	valid?	physical page #
00	1	010
01	1	111
10	0	000
11	1	000

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	CB 0B CB 0B
$0x38-B$	DC 0C DC 0C
$0x3C-F$	EC 0C EC 0C

exercise

5-bit virtual addresses, 6-bit physical addresses, 8-byte pages

(VAs) $0x18 = (11000 \rightarrow 000000) 00$; $0x03 = (00011 \rightarrow 010011) 0x4A$;
 $0x0A = (01010 \rightarrow 111010) 0xDC$; $0x13 = (10011 \rightarrow ?)$ fault

page table


virtual page #	valid?	physical page #
00	1	010
01	1	111
10	0	000
11	1	000

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C


physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	CB 0B CB 0B
$0x38-B$	DC 0C DC 0C
$0x3C-F$	EC 0C EC 0C

emacs (two copies)

Emacs (run by user mst3k)




Emacs (run by user xyz4w)

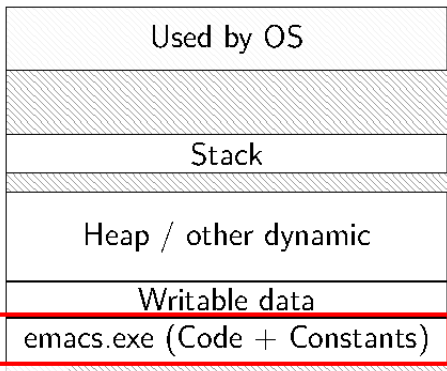


emacs (two copies)

Emacs (run by user mst3k)



Emacs (run by user xyz4w)



same data?

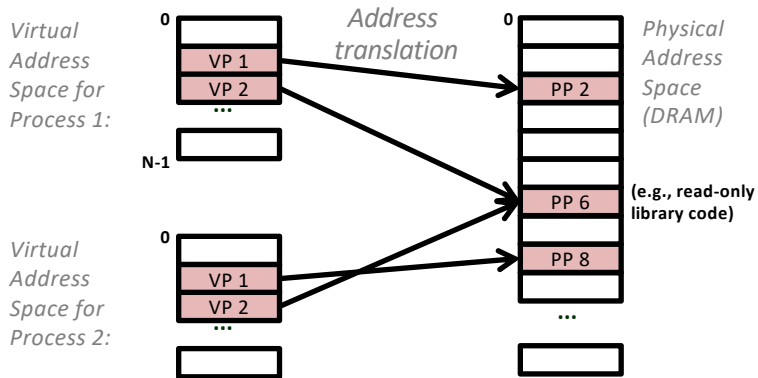
two copies of program

would like to only have one copy of program

what if mst3k's emacs tries to modify its code?

would break process abstraction:

“illusion of own memory”



typical page table entries

solution: same idea as kernel-only bit

page table entry will have more **permissions bits**

can read?

can write?

can execute?

checked by MMU like valid/kernel bit

page table (logically)

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

page tables Summary

program memory = virtual; real memory = physical

each memory divided into fixed-sizes **pages**

each virtual page has a **page table entry**:

- has location of physical page (if any)

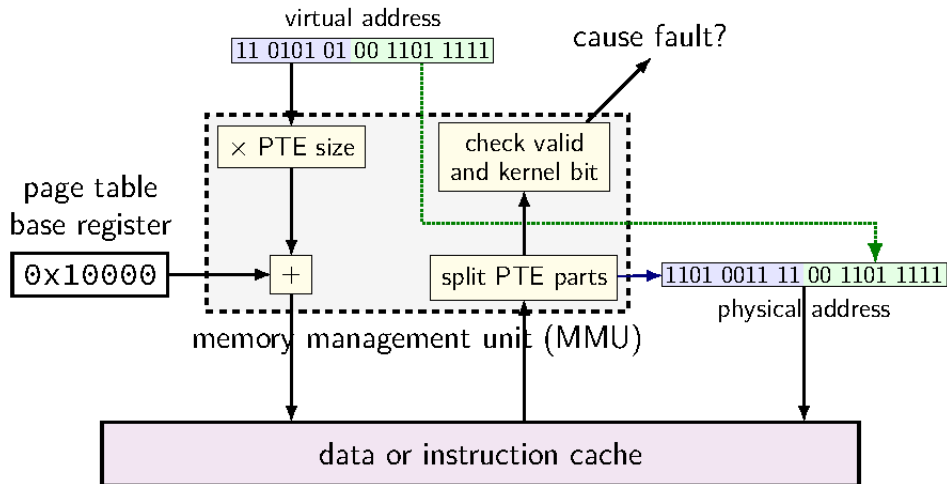
- has valid bit

- has permission bits

all page table entries stored in page table

permission or validity error triggers **exception**

memory access with page table




1-level example

6-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 other bits; page table base register $0x20$; translate virtual address $0x30$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 10 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	CB 0B CB 0B
$0x38-B$	DC 0C DC 0C
$0x3C-F$	EC 0C EC 0C

memory access with page table



1-level example

6-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 other bits; page table base register $0x20$; translate virtual address $0x30$

physical
addresses bytes

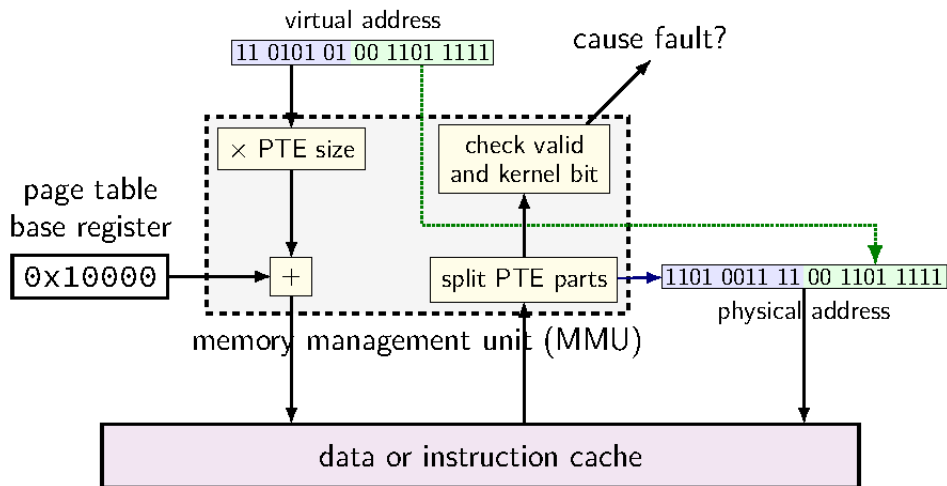
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical
addresses bytes

$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 10 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	CB 0B CB 0B
$0x38-B$	DC 0C DC 0C
$0x3C-F$	EC 0C EC 0C

$0x30 = 11\ 0000$

memory access with page table



1-level example

6-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 other bits; page table base register $0x20$; translate virtual address $0x30$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C


physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 10 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	CB 0B CB 0B
$0x38-B$	DC 0C DC 0C
$0x3C-F$	EC 0C EC 0C

$0x30 = 11\ 0000$

PTE addr:

$0x20 + 6 \times 1 = 0x26$

memory access with page table



1-level example

6-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 other bits; page table base register $0x20$; translate virtual address $0x30$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 10 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	CB 0B CB 0B
$0x38-B$	DC 0C DC 0C
$0x3C-F$	EC 0C EC 0C

$0x30 = 11\ 0000$

PTE addr:

$0x20 + 6 \times 1 = 0x26$

PTE value:

$0x10 = 0001\ 0000$

PPN 000, valid 1

Physical Address 000 000

exercise: 64-bit system

my desktop: 39-bit physical addresses; 48-bit virtual addresses

4096 byte pages

exercise: 64-bit system

my desktop: 39-bit physical addresses; 48-bit virtual addresses

4096 byte pages

top 16 bits of address not used for translation

exercise: 64-bit system

my desktop: 39-bit physical addresses; 48-bit virtual addresses

4096 byte pages

exercise: how many page table entries?

exercise: how large are physical page numbers?

exercise: 64-bit system

my desktop: 39-bit physical addresses; 48-bit virtual addresses

4096 byte pages

exercise: how many page table entries? $2^{48}/2^{12} = 2^{36}$ entries

exercise: how large are physical page numbers? $39 - 12 = 27$ bits

exercise: 64-bit system

my desktop: 39-bit physical addresses; 48-bit virtual addresses

4096 byte pages

exercise: how many page table entries? $2^{48}/2^{12} = 2^{36}$ entries

exercise: how large are physical page numbers? $39 - 12 = 27$ bits

page table entries are **8 bytes** (room for expansion, metadata)

would take up 2^{39} bytes?? (512GB??)


huge page tables

huge virtual address spaces!

impossible to store PTE for every page

how can we save space?

holes



most pages are **invalid**

saving space

basic idea: don't store (most) invalid page table entries

use a data structure other than a flat array

want a map — lookup key (virtual page number), get value (PTE)

options?

saving space

basic idea: don't store (most) invalid page table entries

use a data structure other than a flat array

want a map — lookup key (virtual page number), get value (PTE)

options?

hashtable

actually used by some historical processors
but never common

saving space

basic idea: don't store (most) invalid page table entries

use a data structure other than a flat array

want a map — lookup key (virtual page number), get value (PTE)

options?

hashtable

actually used by some historical processors
but never common

tree data structure

but not quite a search tree

search tree tradeoffs

lookup usually implemented **in hardware**

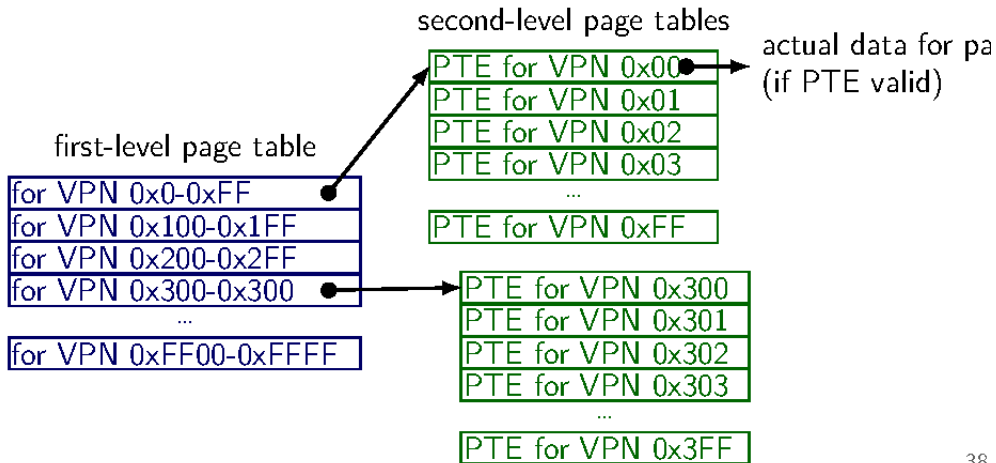
lookup should be simple

lookup should not involve many memory accesses

doing two memory accesses is already very slow

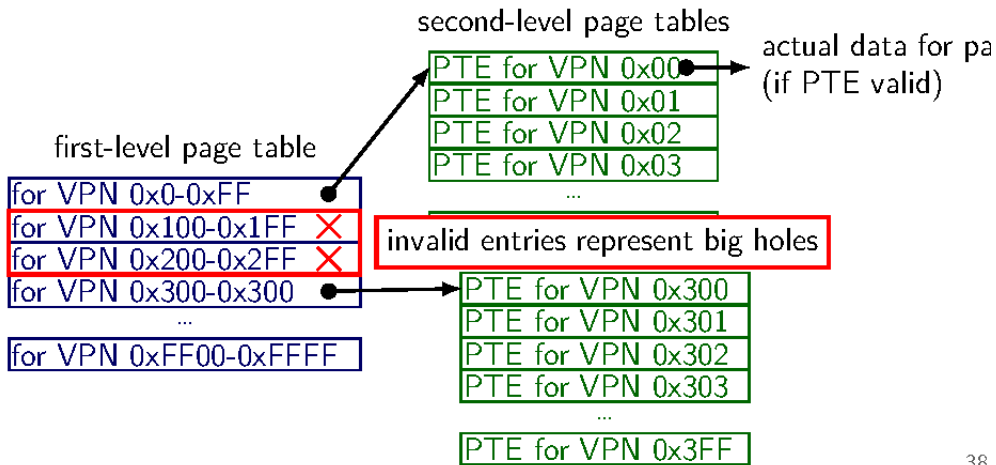
two-level page tables

two-level page table for 65536 pages (16-bit VPN)




two-level page tables

two-level page table for 65536 pages (16-bit VPN)



holes



most pages are **invalid**

two-level page tables

two-level page table for 65536 pages (16-bit VPN)

first-level page table				
VPN range	valid	kernel	write	physical page # (of next page table)
0x0000-0x00FF	1	0	1	0x22343
0x0100-0x01FF	0	0	1	0x00000
0x0200-0x02FF	0	0	0	0x00000
0x0300-0x03FF	1	1	0	0x33454
0x0400-0x04FF	1	1	0	0xFF043
...
0xFF00-0xFFFF	1	1	0	0xFF045

first-level page table
for VPN 0x0-0xFF
for VPN 0x100-0x1FF
for VPN 0x200-0x2FF
for VPN 0x300-0x3FF
...
for VPN 0xFF00-0xFFFF

PTE for VPN 0x303

...

PTE for VPN 0x3FF

two-level page tables

two-level page table for 65536 pages (16-bit VPN)

first-level page table				
VPN range	valid	kernel	write	physical page # (of next page table)
0x0000-0x00FF	1	0	1	0x22343
0x0100-0x01FF	0	0	1	0x00000
0x0200-0x02FF	0	0	0	0x00000
0x0300-0x03FF	1	1	0	0x33454
0x0400-0x04FF	1	1	0	0xFF043
...
0xFF00-0xFFFF	1	1	0	0xFF045

first-level page table
for VPN 0x0-0xFF
for VPN 0x100-0x1FF
for VPN 0x200-0x2FF
for VPN 0x300-0x3FF
...
for VPN 0xFF00-0xFFFF

PTE for VPN 0x303

...

PTE for VPN 0x3FF

two-level page tables

two-level page table for 65536 pages (16-bit VPN)

		first-level page table				
		VPN range	valid	kernel	write	physical page # (of next page table)
first-level page for VPN 0x0-0xFF for VPN 0x100-0x1FF for VPN 0x200-0x2FF for VPN 0x300-0x3FF ... for VPN 0xFF00-0xFF0F	0x0000-0x00FF	1	0	1	0x22343	
	0x0100-0x01FF	0	0	1	0x00000	
	0x0200-0x02FF	0	0	0	0x00000	
	0x0300-0x03FF	1	1	0	0x33454	
	0x0400-0x04FF	1	1	0	0xFF043	
	
	0xFF00-0xFFFF	1	1	0	0xFF045	

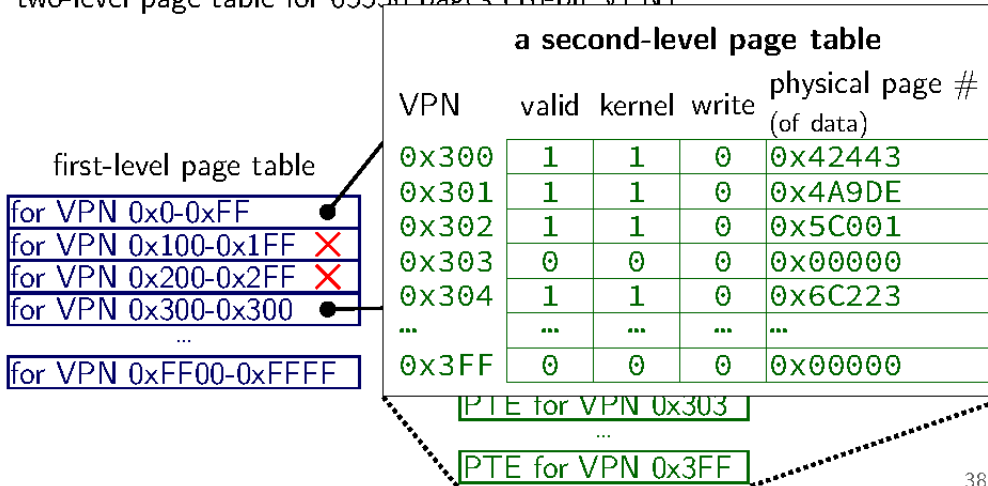
PTE for VPN 0x303

...

PTE for VPN 0x3FF

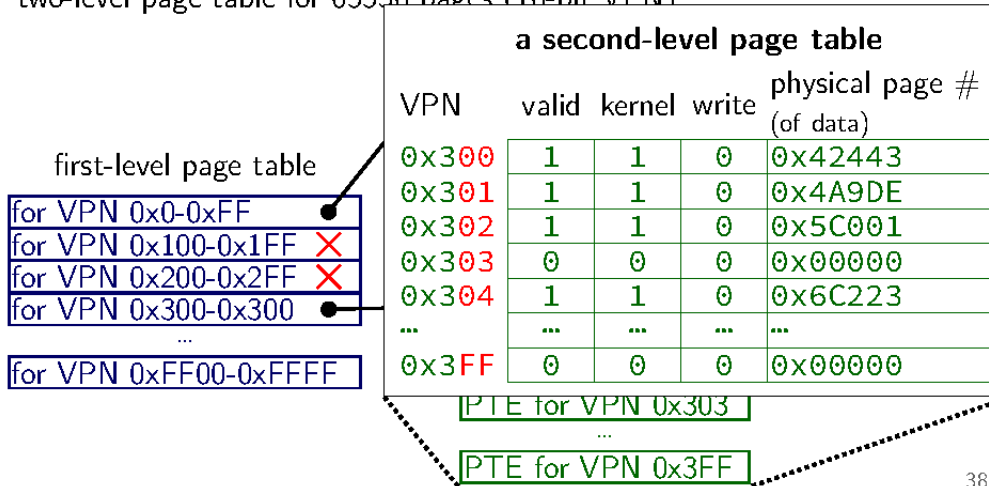
two-level page tables

two-level page table for 65536 pages (16-bit VPN)



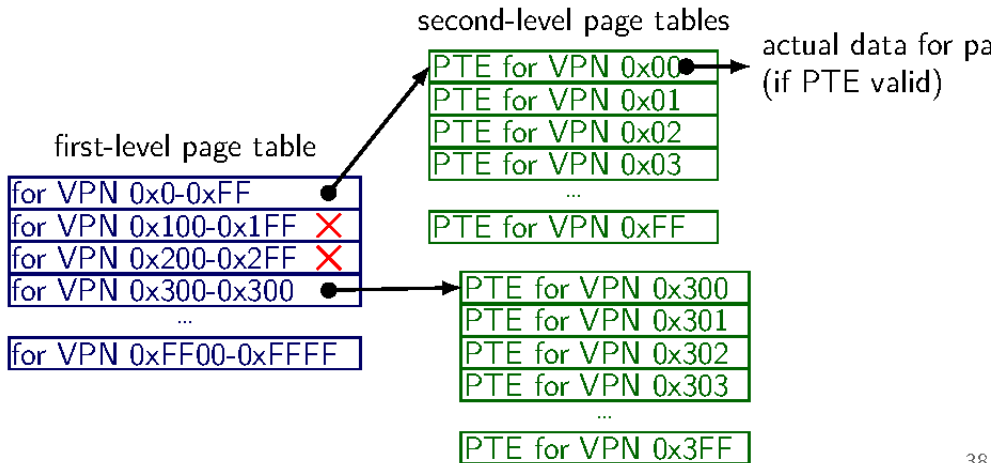
two-level page tables

two-level page table for 65536 pages (16-bit VPN)



two-level page tables

two-level page table for 65536 pages (16-bit VPN)




two-level page table lookup

virtual address

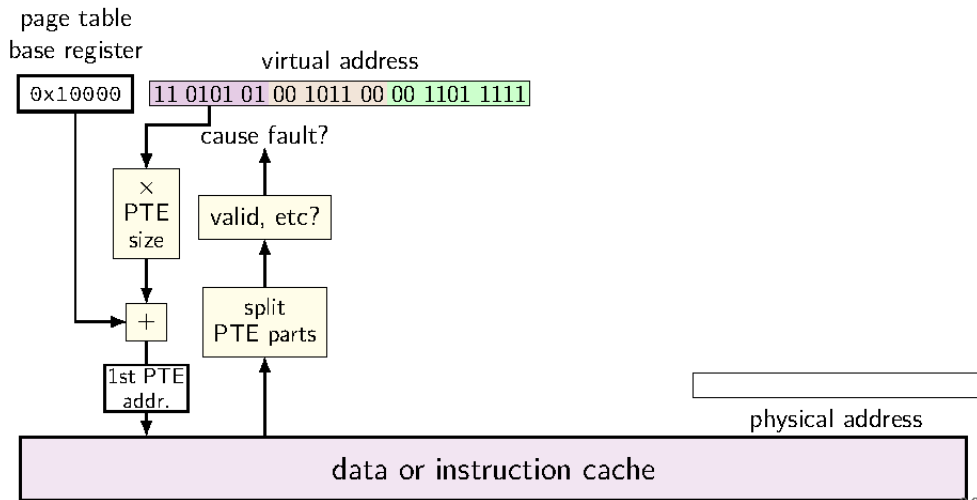
11 0101 01 00 1011 00 00 1101 1111

VPN — split into two parts (one per level)

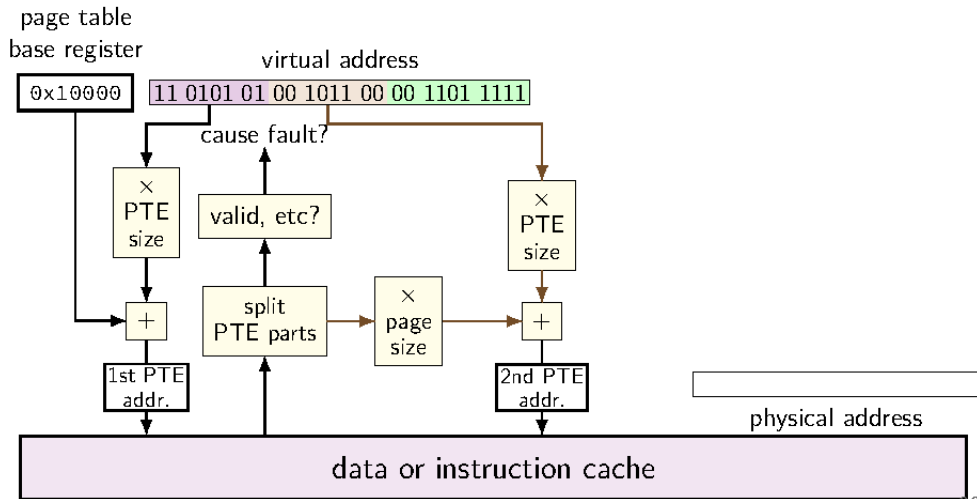
two-level page table lookup



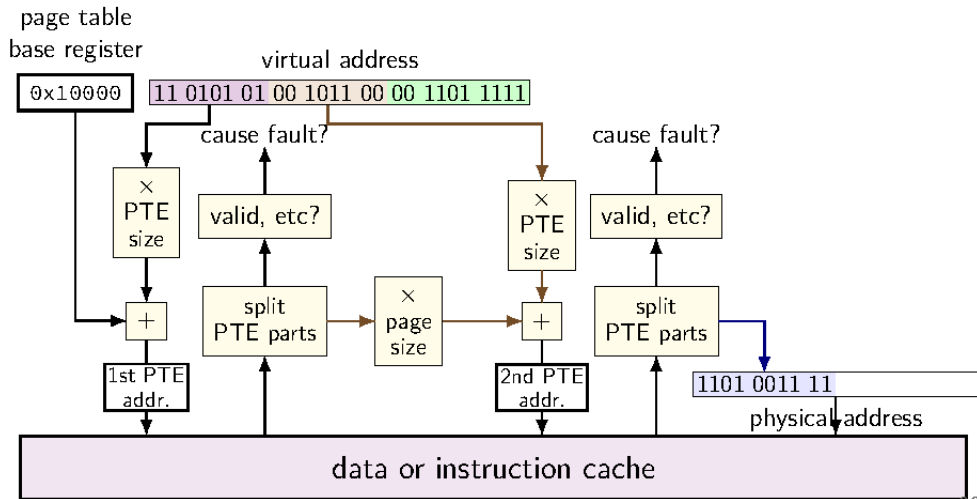
two-level page table lookup



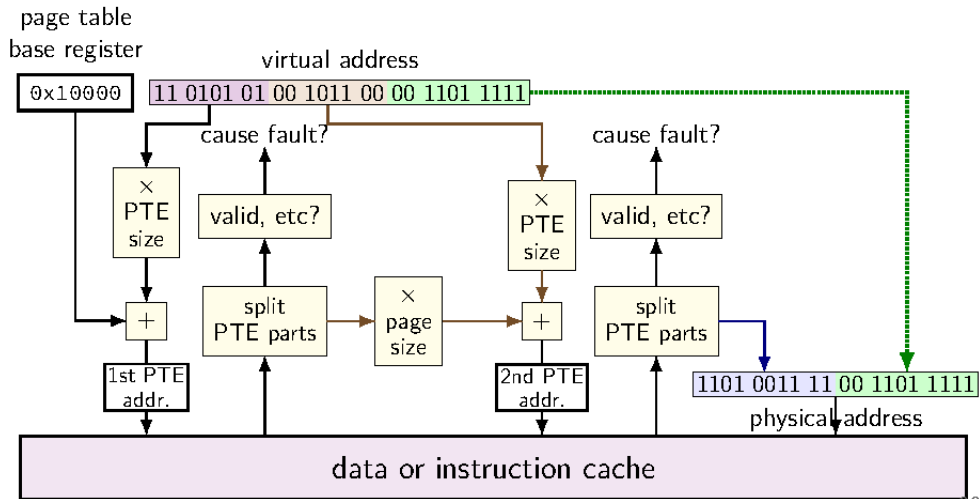
two-level page table lookup



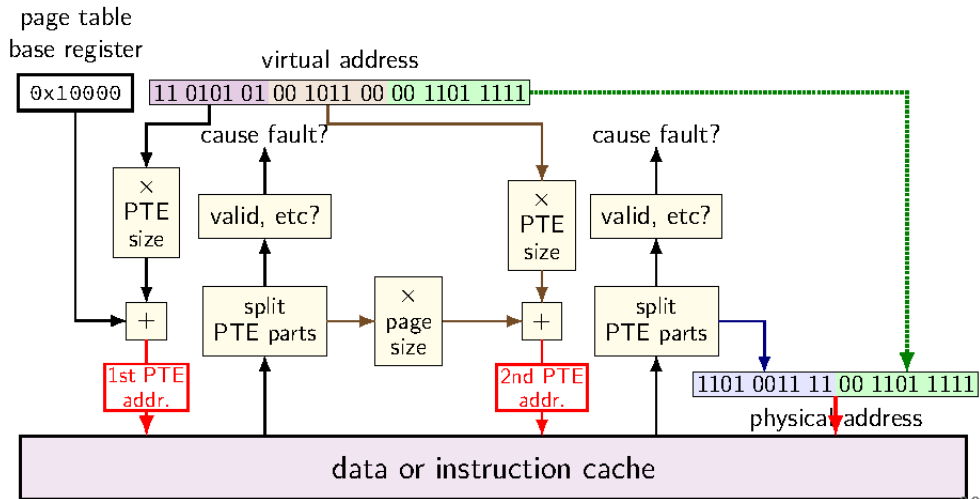
two-level page table lookup



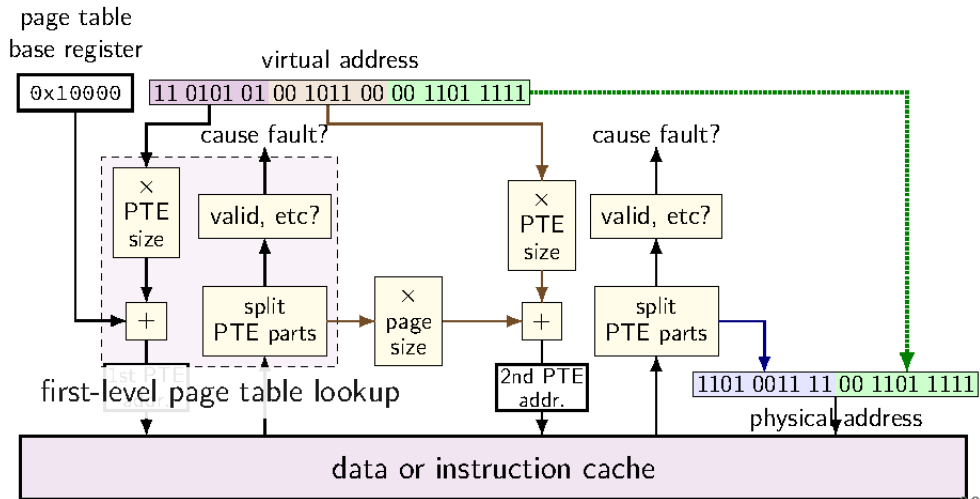
two-level page table lookup



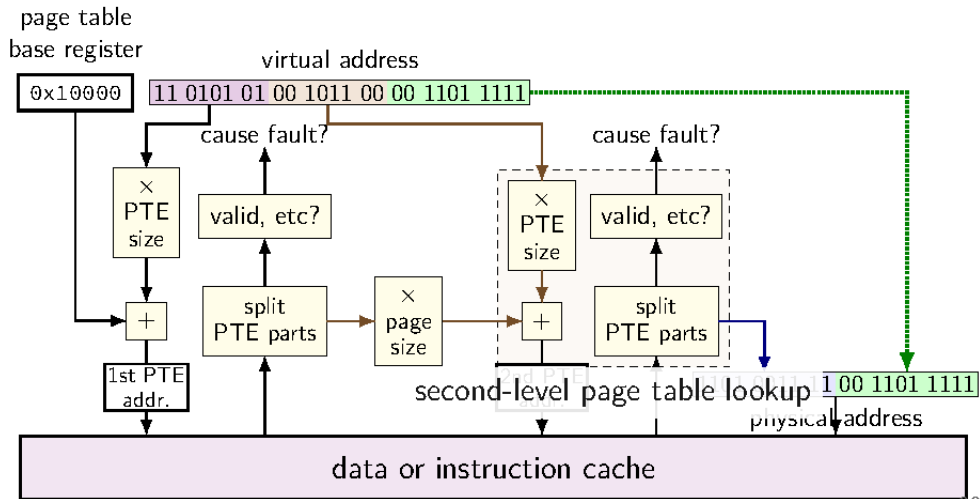
two-level page table lookup



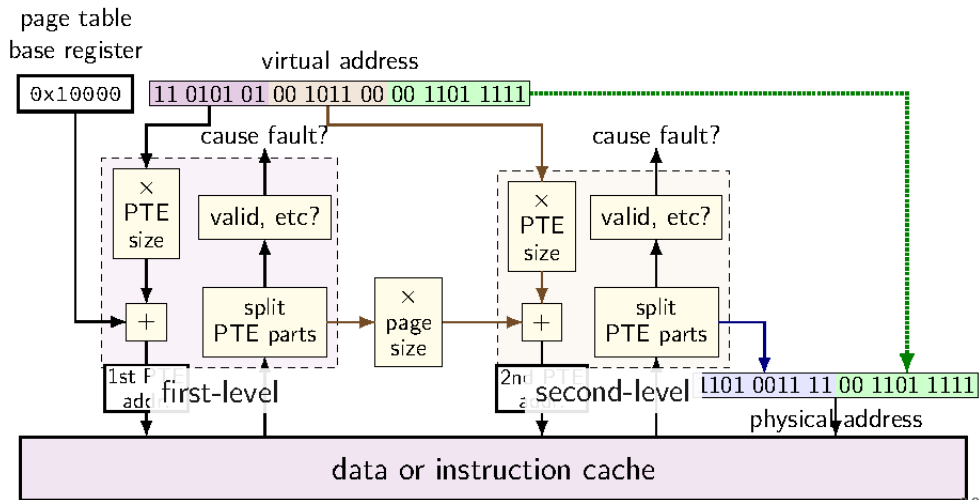
two-level page table lookup



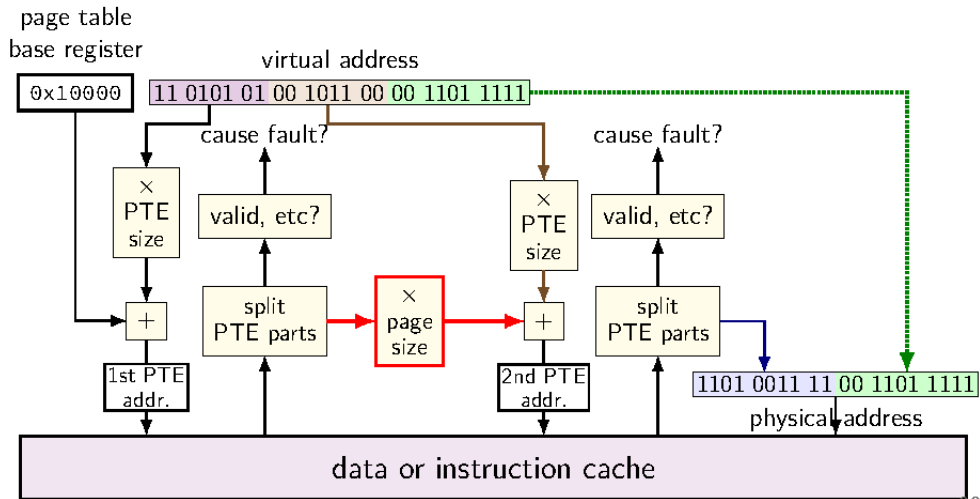
two-level page table lookup




two-level page table lookup



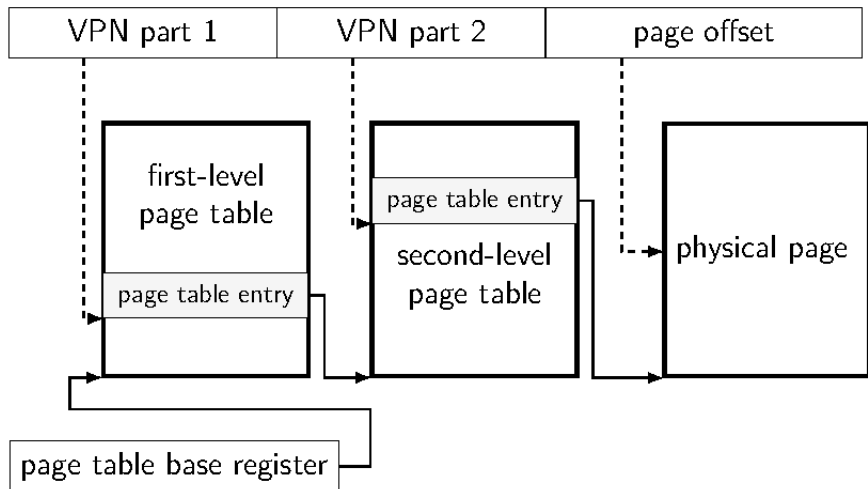
two-level page table lookup



two-level page table lookup



another view



multi-level page tables

VPN split into pieces for each level of page table

top levels: page table entries point to next page table
usually using physical page number of next page table

bottom level: page table entry points to destination page

validity and permission checks at **each level**

note on VPN splitting

textbook labels it 'VPN 1' and 'VPN 2' and so on

these are **parts of the virtual page number**

(there are not multiple VPNs)

splitting addresses for levels

x86-32

32-bit physical address; 32-bit virtual address

2^{12} byte page size

2-levels of page tables; each page table is one page

4 byte page table entries

how is address 0x12345678 split up?

splitting addresses for levels

x86-32

32-bit physical address; 32-bit virtual address

2^{12} byte page size

12-bit page offset

2-levels of page tables; each page table is one page

4 byte page table entries

how is address 0x12345678 split up?

splitting addresses for levels

x86-32

32-bit physical address; 32-bit virtual address

2^{12} byte page size

12-bit page offset

2-levels of page tables; each page table is one page

4 byte page table entries

$2^{12}/4 = 2^{10}$ PTEs/page table; *10-bit VPN parts*

how is address 0x12345678 split up?

splitting addresses for levels

x86-32

32-bit physical address; 32-bit virtual address

2^{12} byte page size

12-bit page offset

2-levels of page tables; each page table is one page

4 byte page table entries

$2^{12}/4 = 2^{10}$ PTEs/page table; *10-bit VPN parts*

how is address 0x12345678 split up?

10-bit VPN part 1: 0001 0010 00 (0x48);

10-bit VPN part 2: 11 0100 0101 (0x345);

12-bit page offset: 0x678

2-level example

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x20$; translate virtual address $0x131$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D0 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

2-level splitting

9-bit virtual address

6-bit physical address

8-byte pages \rightarrow 3-bit page offset (bottom bits)

9-bit VA: 6 bit VPN + 3 bit PO

6-bit PA: 6 bit PPN + 3 bit PO

8 entry page tables \rightarrow 3-bit VPN parts

9-bit VA: 3 bit VPN part 1; 3 bit VPN part 2

2-level example

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE


page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x20$; translate virtual address $0x131$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D0 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

two-level page table lookup



2-level example

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x20$; translate virtual address $0x131$

physical bytes
addresses

$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C


physical bytes
addresses

$0x20-3$	D0 D1 D2 D3
$0x24-7$	D0 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x131 = 1\ 0011\ 0001$

$0x20 + 4 \times 1 = 0x24$

two-level page table lookup



2-level example

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x20$; translate virtual address $0x131$

physical bytes
addresses

$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical bytes
addresses

$0x20-3$	D0 D1 D2 D3
$0x24-7$	D0 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x131 = 1\ 0011\ 0001$


$0x20 + 4 \times 1 = 0x24$

PTE 1 value:

$0xD0 = 1101\ 0000$

PPN 110, valid 1

two-level page table lookup



2-level example

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x20$; translate virtual address $0x131$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D0 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x131 = 1\ 0011\ 0001$

$0x20 + 4 \times 1 = 0x24$

PTE 1 value:

$0xD0 = 1101\ 0000$

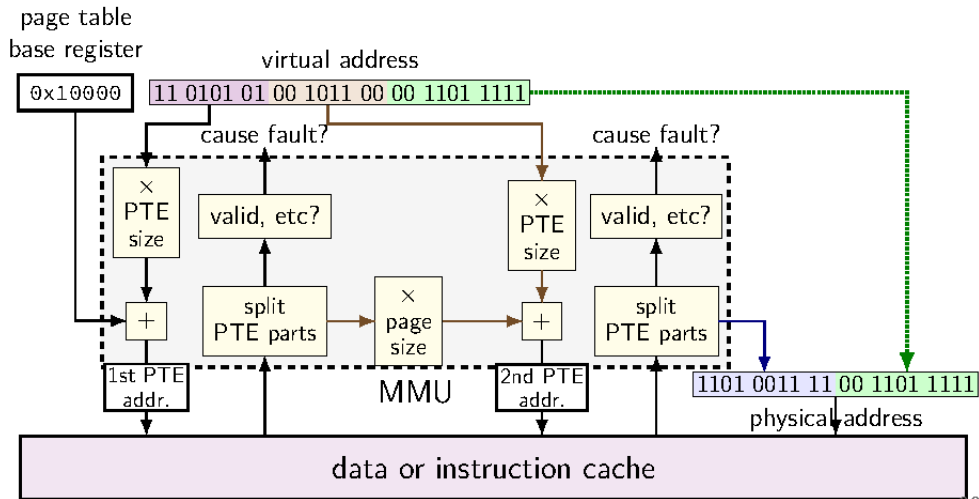
PPN 110, valid 1

PTE 2 addr:

$110\ 000 + 110 = 0x36$

PTE 2 value: $0xDB$

two-level page table lookup



2-level example

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x20$; translate virtual address $0x131$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x131 = 1\ 0011\ 0001$

$0x20 + 4 \times 1 = 0x24$

PTE 1 value:

$0xD4 = 1101\ 0100$

PPN 110, valid 1

PTE 2 addr:

$110\ 000 + 110 = 0x36$

PTE 2 value: $0xDB$

PPN 110; valid 1

$M[110\ 001\ (0x31)] = 0x0A$

2-level exercise (1)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused;

page table base register $0x08$; translate virtual address $0x0FB$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

2-level exercise (1)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused;

page table base register $0x08$; translate virtual address $0x0FB$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x0FB = 011\ 111\ 011$

PTE 1 addr:

$0x08 + 3 \times 1 = 0x0B$

PTE 1: $0xBB$ at $0x0B$

PTE 1: PPN 101 (5) valid 1

PTE 2 addr:

$101\ 000 + 111 = 0x2F$

PTE 2: $0xF0$ at $0x2F$

PTE 2: PPN 111 (7) valid 1

$111\ 011 = 0x3B \rightarrow$

$0x0C$

2-level exercise (1)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused;

page table base register $0x08$; translate virtual address $0x0FB$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x0FB = 011\ 111\ 011$

PTE 1 addr:

$0x08 + 3 \times 1 = 0x0B$

PTE 1: **$0xBB$** at $0x0B$

PTE 1: PPN 101 (5) valid 1

PTE 2 addr:

$101\ 000 + 111 = 0x2F$

PTE 2: $0xF0$ at $0x2F$

PTE 2: PPN 111 (7) valid 1

$111\ 011 = 0x3B \rightarrow$
 $0x0C$

2-level exercise (1)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused;

page table base register $0x08$; translate virtual address $0x0FB$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x0FB = 011\ 111\ 011$

PTE 1 addr:

$0x08 + 3 \times 1 = 0x0B$

PTE 1: $0xBB$ at $0x0B$

PTE 1: PPN 101 (5) valid 1

PTE 2 addr:

$101\ 000 + 111 = 0x2F$

PTE 2: **$0xF0$** at $0x2F$

PTE 2: PPN 111 (7) valid 1

$111\ 011 = 0x3B \rightarrow$
 $0x0C$

2-level exercise (1)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused;

page table base register $0x08$; translate virtual address $0x0FB$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x0FB = 011\ 111\ 011$

PTE 1 addr:

$0x08 + 3 \times 1 = 0x0B$

PTE 1: $0xBB$ at $0x0B$

PTE 1: PPN 101 (5) valid 1

PTE 2 addr:

$101\ 000 + 111 = 0x2F$

PTE 2: $0xF0$ at $0x2F$

PTE 2: PPN 111 (7) valid 1

$111\ 011 = 0x3B \rightarrow$
 $0x0C$

2-level exercise (2)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x08$; translate virtual address $0x00B$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

2-level exercise (2)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x08$; translate virtual address $0x00B$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x00B = 000\ 001\ 011$

PTE 1: $0x88$ at $0x08$

PTE 1: PPN 100 (5) valid 0
page fault!

2-level exercise (2)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x08$; translate virtual address $0x00B$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x00B = 000\ 001\ 011$

PTE 1: $0x88$ at $0x08$

PTE 1: PPN 100 (5) valid 0
page fault!

2-level exercise (3)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x08$; translate virtual address $0x1CB$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

2-level exercise (3)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x08$; translate virtual address $0x1CB$

physical bytes
addresses

$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical bytes
addresses

$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x1CB = 111\ 001\ 011$

PTE 1: $0xFF$ at $0x0F$

PTE 1: PPN 111 (7) valid 1

PTE 2: $0x0C$ at $0x39$

PTE 2: PPN 000 (0) valid 0

page fault!

2-level exercise (3)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x08$; translate virtual address $0x1CB$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x1CB = 111\ 001\ 011$

PTE 1: $0xFF$ at $0x0F$

PTE 1: PPN 111 (7) valid 1

PTE 2: $0x0C$ at $0x39$

PTE 2: PPN 000 (0) valid 0

page fault!

2-level exercise (3)

9-bit virtual addresses, 6-bit physical; 8 byte pages, 1 byte PTE

page tables 1 page; PTE: 3 bit PPN (MSB), 1 valid bit, 4 unused

page table base register $0x08$; translate virtual address $0x1CB$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	1C 2C 3C 4C

physical addresses	bytes
$0x20-3$	D0 D1 D2 D3
$0x24-7$	D4 D5 D6 D7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x1CB = 111\ 001\ 011$

PTE 1: $0xFF$ at $0x0F$

PTE 1: PPN 111 (7) valid 1

PTE 2: $0x0C$ at $0x39$

PTE 2: PPN 000 (0) valid 0

page fault!

2-level exercise (4)

10-bit virtual addresses, 6-bit physical; 16 byte pages, 2 byte PTE

page tables 1 page; PTE: 2 bit PPN (MSB of first byte), 1 valid bit, rest unused

page table base register $0x10$; translate virtual address $0x376$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	AC BC DC EC

physical addresses	bytes
$0x20-3$	D0 E1 D2 D3
$0x24-7$	D4 E5 D6 E7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

2-level exercise (4)

10-bit virtual addresses, 6-bit physical; 16 byte pages, 2 byte PTE

page tables 1 page; PTE: 2 bit PPN (MSB of first byte), 1 valid bit, rest unused

page table base register $0x10$; translate virtual address $0x376$

physical bytes
addresses

$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	AC BC DC EC

physical bytes
addresses

$0x20-3$	D0 E1 D2 D3
$0x24-7$	D4 E5 D6 E7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x376 = 110\ 111\ 0110$

PTE 1: $0x10 + \mathbf{6} \times \mathbf{2} = 0x1C$:
AC BC

PTE 1: PPN 10 valid 1

PTE 2: $0x20 + \mathbf{7} \times \mathbf{2} = 0x2E$:
EF F0

PTE 2: PPN 11 valid 1

11 0110 = $0x36 \rightarrow$
DB

2-level exercise (4)

10-bit virtual addresses, 6-bit physical; 16 byte pages, 2 byte PTE

page tables 1 page; PTE: 2 bit PPN (MSB of first byte), 1 valid bit, rest unused

page table base register $0x10$; translate virtual address $0x376$

physical bytes
addresses

$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	AC BC DC EC

physical bytes
addresses

$0x20-3$	D0 E1 D2 D3
$0x24-7$	D4 E5 D6 E7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x376 = 110\ 111\ 0110$

PTE 1: $0x10 + 6 \times 2 = 0x1C$:
AC BC

PTE 1: PPN 10 valid 1

PTE 2: $0x20 + 7 \times 2 = 0x2E$:
EF F0

PTE 2: PPN 11 valid 1

$11\ 0110 = 0x36 \rightarrow$
DB

2-level exercise (4)

10-bit virtual addresses, 6-bit physical; 16 byte pages, 2 byte PTE

page tables 1 page; PTE: 2 bit PPN (MSB of first byte), 1 valid bit, rest unused

page table base register $0x10$; translate virtual address $0x376$

physical
addresses bytes

$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	AC BC DC EC

physical
addresses bytes

$0x20-3$	D0 E1 D2 D3
$0x24-7$	D4 E5 D6 E7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x376 = 110\ 111\ 0110$

PTE 1: $0x10 + \mathbf{6} \times \mathbf{2} = 0x1C$:

AC BC

PTE 1: PPN 10 valid 1

PTE 2: $0x20 + \mathbf{7} \times \mathbf{2} = 0x2E$:

EF F0

PTE 2: PPN 11 valid 1

11 0110 = $0x36 \rightarrow$

DB

2-level exercise (4)

10-bit virtual addresses, 6-bit physical; 16 byte pages, 2 byte PTE

page tables 1 page; PTE: 2 bit PPN (MSB of first byte), 1 valid bit, rest unused

page table base register $0x10$; translate virtual address $0x376$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	AC BC DC EC

physical addresses	bytes
$0x20-3$	D0 E1 D2 D3
$0x24-7$	D4 E5 D6 E7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x376 = 110 \text{ } 111 \text{ } 0110$
PTE 1: $0x10 + 6 \times 2 = 0x1C$:
AC BC
PTE 1: PPN 10 valid 1
PTE 2: $0x20 + 7 \times 2 = 0x2E$:
EF F0
PTE 2: PPN 11 valid 1
 $11 \text{ } 0110 = 0x36 \rightarrow$
DB

2-level exercise (4)

10-bit virtual addresses, 6-bit physical; 16 byte pages, 2 byte PTE

page tables 1 page; PTE: 2 bit PPN (MSB of first byte), 1 valid bit, rest unused

page table base register $0x10$; translate virtual address $0x376$

physical bytes
addresses

$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	AC BC DC EC

physical bytes
addresses

$0x20-3$	D0 E1 D2 D3
$0x24-7$	D4 E5 D6 E7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF FO
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x376 = 110\ 111\ 0110$

PTE 1: $0x10 + \mathbf{6} \times \mathbf{2} = 0x1C$:
AC BC

PTE 1: PPN 10 valid 1

PTE 2: $0x20 + \mathbf{7} \times \mathbf{2} = 0x2E$:
EF FO

PTE 2: PPN 11 valid 1

11 0110 = $0x36 \rightarrow$
DB

2-level exercise (4)

10-bit virtual addresses, 6-bit physical; 16 byte pages, 2 byte PTE

page tables 1 page; PTE: 2 bit PPN (MSB of first byte), 1 valid bit, rest unused

page table base register $0x10$; translate virtual address $0x376$

physical addresses	bytes
$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	AC BC DC EC

physical addresses	bytes
$0x20-3$	D0 E1 D2 D3
$0x24-7$	D4 E5 D6 E7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x376 = 110\ 111\ 0110$
PTE 1: $0x10 + \mathbf{6} \times \mathbf{2} = 0x1C$:
AC BC
PTE 1: PPN 10 valid 1
PTE 2: $0x20 + \mathbf{7} \times \mathbf{2} = 0x2E$:
EF F0
PTE 2: PPN 11 valid 1
 $11\ 0110 = 0x36 \rightarrow$
DB

2-level exercise (4)

10-bit virtual addresses, 6-bit physical; 16 byte pages, 2 byte PTE

page tables 1 page; PTE: 2 bit PPN (MSB of first byte), 1 valid bit, rest unused

page table base register $0x10$; translate virtual address $0x376$

physical bytes
addresses

$0x00-3$	00 11 22 33
$0x04-7$	44 55 66 77
$0x08-B$	88 99 AA BB
$0x0C-F$	CC DD EE FF
$0x10-3$	1A 2A 3A 4A
$0x14-7$	1B 2B 3B 4B
$0x18-B$	1C 2C 3C 4C
$0x1C-F$	AC BC DC EC

physical bytes
addresses

$0x20-3$	D0 E1 D2 D3
$0x24-7$	D4 E5 D6 E7
$0x28-B$	89 9A AB BC
$0x2C-F$	CD DE EF F0
$0x30-3$	BA 0A BA 0A
$0x34-7$	DB 0B DB 0B
$0x38-B$	EC 0C EC 0C
$0x3C-F$	FC 0C FC 0C

$0x376 = 110\ 111\ 0110$

PTE 1: $0x10 + \mathbf{6} \times \mathbf{2} = 0x1C$:
AC BC

PTE 1: PPN 10 valid 1

PTE 2: $0x20 + \mathbf{7} \times \mathbf{2} = 0x2E$:
EF F0

PTE 2: PPN 11 valid 1

11 0110 = $0x36 \rightarrow$

DB

current x86-64 page tables

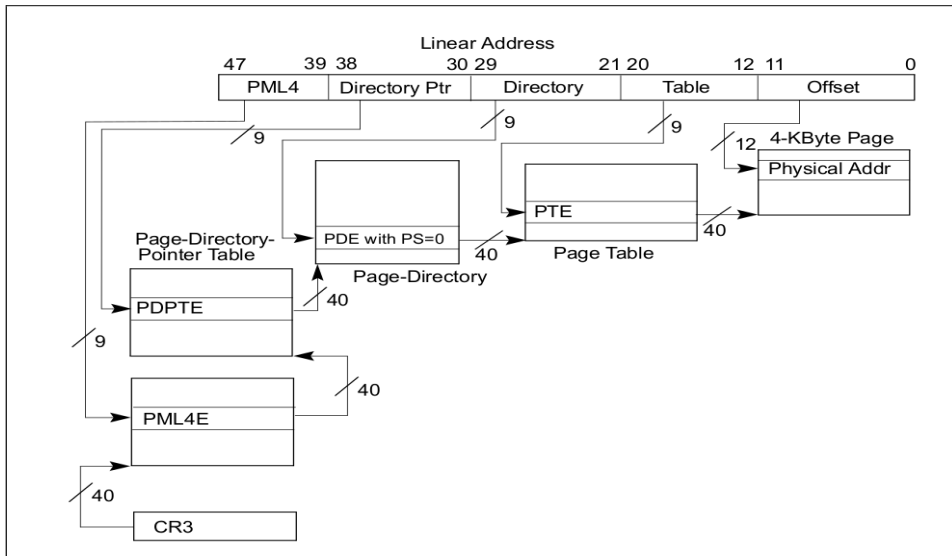



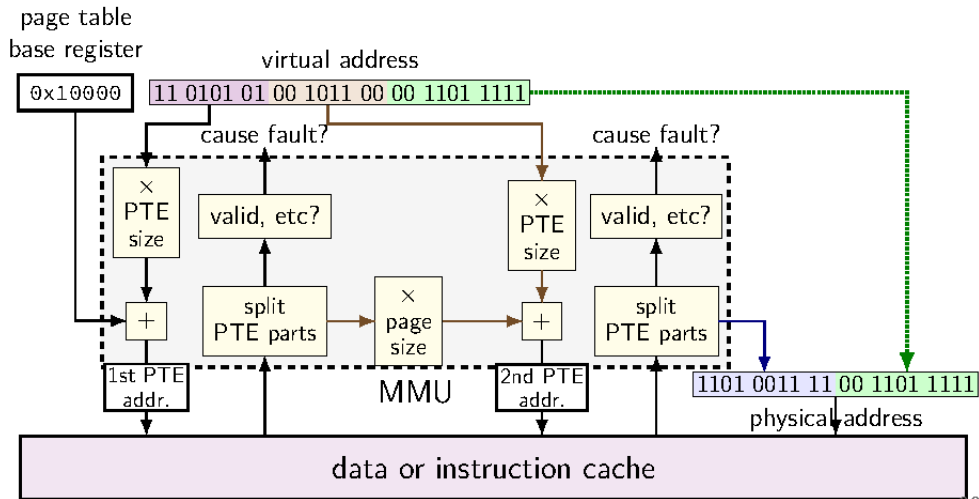
Figure 4-8. Linear-Address Translation to a 4-KByte Page using 4-Level Paging

two-level page tables

two-level page table for 65536 pages (16-bit VPN)



two-level page table lookup



current x86-64 page tables

4-level page table

512 PTEs of 8 bytes each for each page table

choice: exactly one page per page table

allows OS to allocate new page table space in one page units

page table space exercise (1)

4-level page table

512 PTEs of 8 bytes each for each page table

suppose a process has exactly one page allocated

how much space for page tables?

page table space exercise (1)

4-level page table


512 PTEs of 8 bytes each for each page table

suppose a process has exactly one page allocated

how much space for page tables?

1 page at each level (4KB each) exactly

page table space exercise (1)



4 page tables at 1 page/page table
plus 1 page of data
5 pages total (Data)

page table space exercise (2)

4-level page table

512 PTEs of 8 bytes each for each page table

suppose a process has exactly two pages allocated:

one at address `0x0`, one at address `0x200000000000`

how much space for page tables?

page table space exercise (2)

4-level page table

512 PTEs of 8 bytes each for each page table

suppose a process has exactly two pages allocated:

one at address `0x0`, one at address `0x200000000000`

how much space for page tables?


1 shared first-level PT, with two valid entries

two second-level PTs, each with one valid entry

two third-level PTs, each with one valid entry

two fourth-level PTs, each with one valid entry

page table space exercise (2)



page table space exercise (3)

4-level page table; each PT: 512 PTEs of 8 bytes

suppose a process has 100 pages of stack, 100 pages of code+constants (contiguous)

stack and code+constants far apart

how much space for page tables?

page table space exercise (3)

4-level page table; each PT: 512 PTEs of 8 bytes

suppose a process has 100 pages of stack, 100 pages of code+constants (contiguous)

stack and code+constants far apart

how much space for page tables? — minimum:


1 shared first-level PT, with two valid entries

two second-level PT, each with one valid entry

two third-level PT, each with one valid entry

two fourth-level PT, each with 100 valid entries

page table space exercise (3)



page table space exercise (3)

4-level page table; each PT: 512 PTEs of 8 bytes

suppose a process has 100 pages of stack, 100 pages of code+constants (contiguous)

how much space for page tables?

page table space exercise (3)

4-level page table; each PT: 512 PTEs of 8 bytes

suppose a process has 100 pages of stack, 100 pages of code+constants (contiguous)

how much space for page tables? — maximum:

1 shared first-level PT, with four valid entries

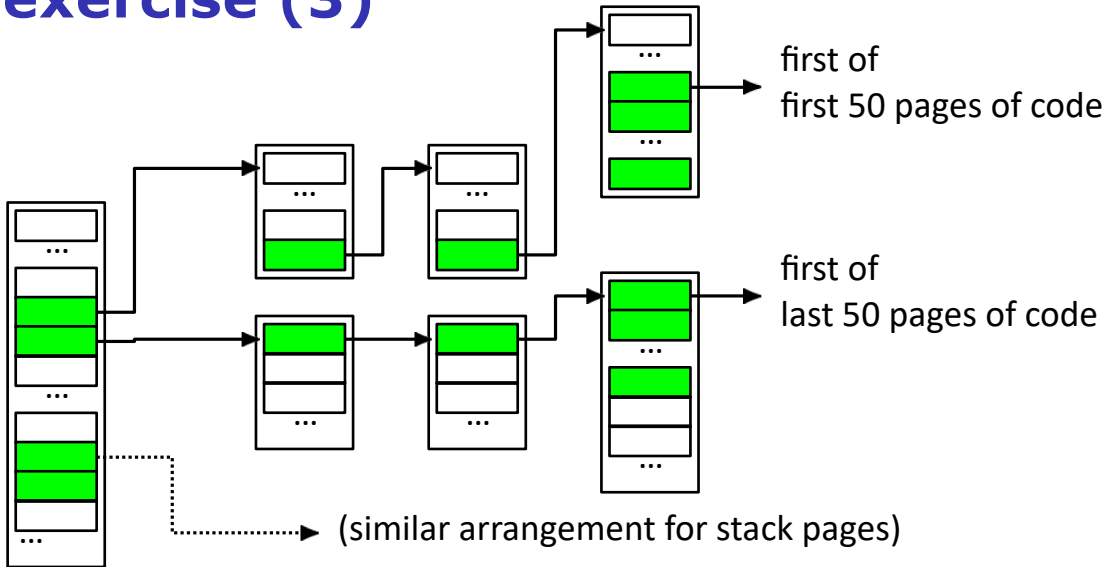
four second-level PT, each with one valid entry

two for stack, two for code+constants

four third-level PT, each with one valid entry

four fourth-level PT, each with 50 valid entries

page table space exercise (3)



page table space exercise (4)

4-level page table; each PT: 512 PTEs of 8 bytes

suppose a process has 200 pages, randomly distributed in PT

about how much space for page tables?

page table space exercise (4)

4-level page table; each PT: 512 PTEs of 8 bytes

suppose a process has 200 pages, randomly distributed in PT

about how much space for page tables?

about 165 ($\pm \sim 8$) entries in first-level PT

(some pages randomly share first-level PT entries)

about 165 second-level PTs, 200 third-level, 200 fourth-level

a bit less than 600 page tables — almost 2400 KB

cache accesses and multi-level PTs

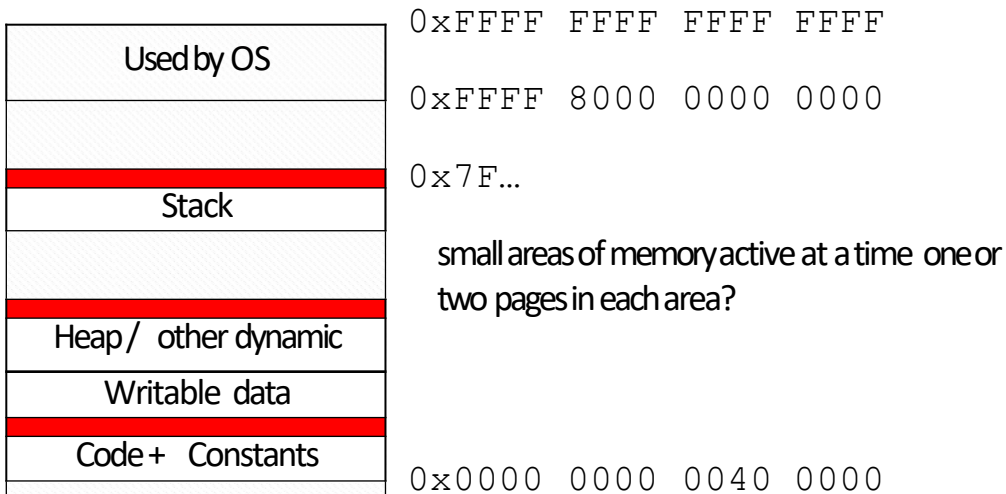
four-level page tables — four cache accesses per memory

access L1 cache hits — typically a couple cycles each?

so add 8 cycles to each memory access?

not acceptable

program memory active sets



page table entries and locality

page table entries have **excellent temporal**

locality typically one or two pages of the stack

active typically one or two pages of code active

typically one or two pages of heap/globals active

each page contains **whole functions**, arrays, stack frames, etc.

page table entries and locality

page table entries have **excellent temporal**

locality typically one or two pages of the stack

active typically one or two pages of code active

typically one or two pages of heap/globals active

each page contains **whole functions**, arrays, stack frames, etc.

needed page table entries are **very small**

cache translated addresses

called a **TLB** (translation lookaside buffer)

small set-associative hardware cache in MMU

maps virtual page numbers to physical page numbers **contains complete page table entries for small number of pages**

L1 cache physical	TLB
addresses	virtual page numbers
bytes from memory	page table entries
tens of bytes per block	one page table entry per block
usually thousands of blocks	usually tens of entries

cache translated addresses

called a **TLB** (translation lookaside buffer)

small set-associative hardware cache in MMU

maps virtual page numbers to physical page numbers **contains complete page table entries for small number of pages**

L1 cache physical	TLB
addresses	virtual page numbers
bytes from memory	page table entries
tens of bytes per block	one page table entry per block
usually thousands of blocks	usually tens of entries
only caches the page table lookup itself (generally) just entries from the last-level page table	

cache translated addresses

called a **TLB** (translation lookaside buffer)

small set-associative hardware cache in MMU

maps virtual page numbers to physical page numbers **contains complete page table entries for small number of pages**

L1 cache physical	TLB
addresses	virtual page numbers
bytes from memory	page table entries
tens of bytes per block	one page table entry per block
usually thousands of blocks	usually tens of entries

not much spatial locality between page table entries
(they're used for kilobytes of data already)

cache translated addresses

called a **TLB** (translation lookaside buffer)

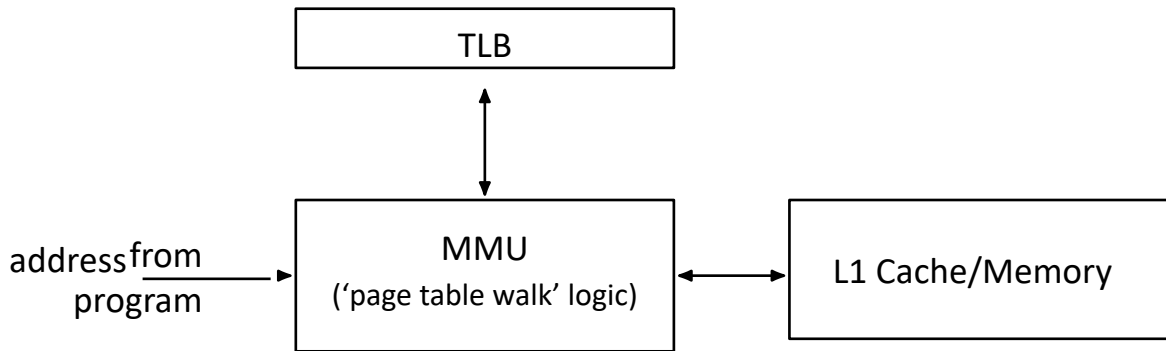
small set-associative hardware cache in MMU

maps virtual page numbers to physical page numbers **contains complete page table entries for small number of pages**

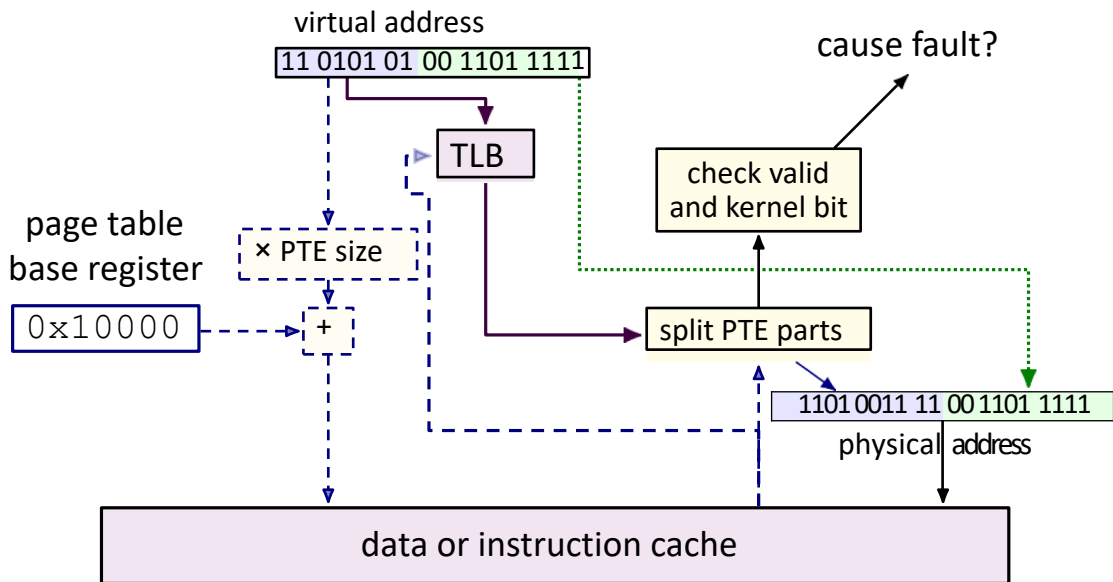
L1 cache physical	TLB
addresses	virtual page numbers
bytes from memory	page table entries
tens of bytes per block	one page table entry per block
usually thousands of blocks	usually tens of entries

few active page table entries at a time
enables highly associative cache designs

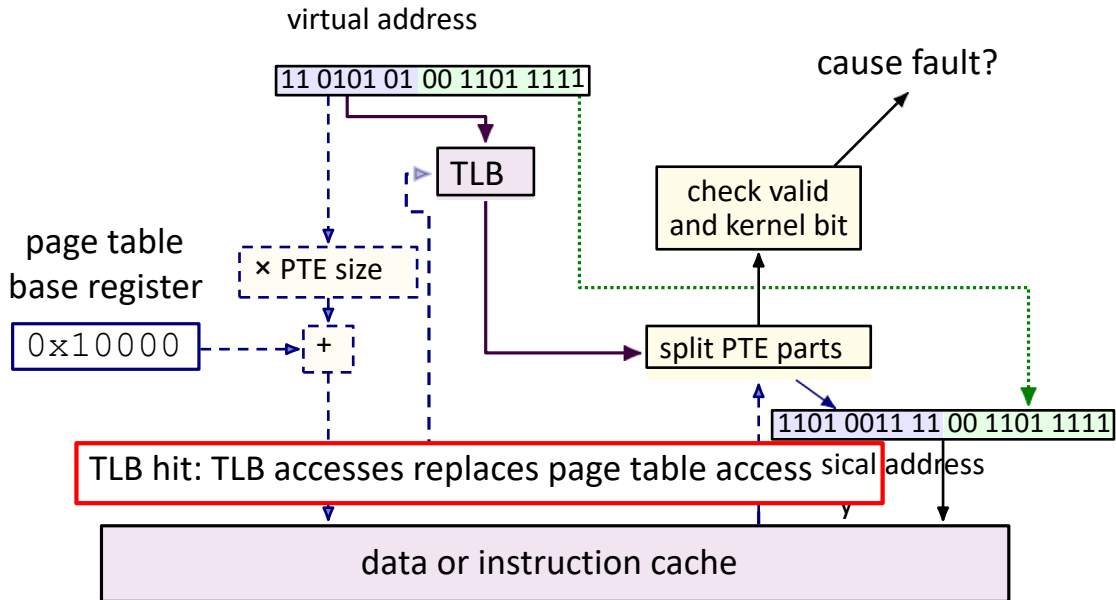
TLB and the MMU (1)



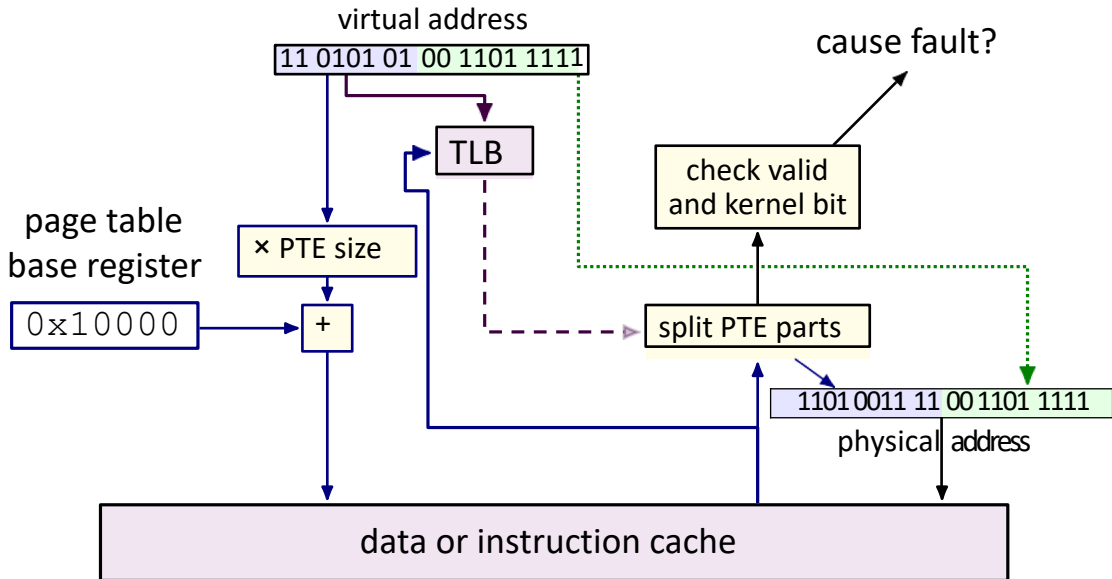
TLB and the MMU (2)



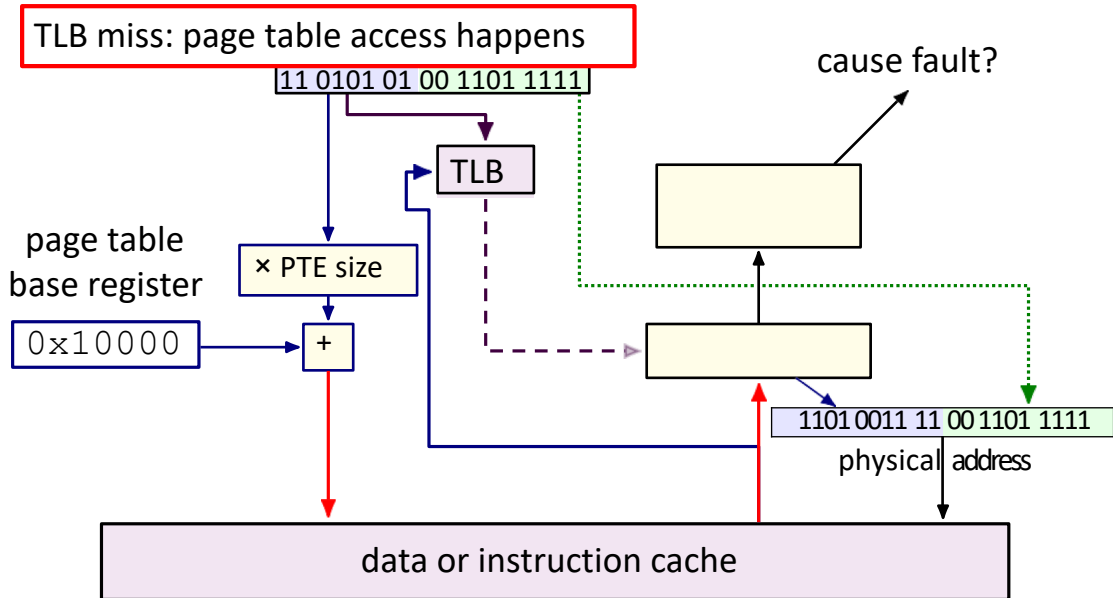
TLB and the MMU (2)



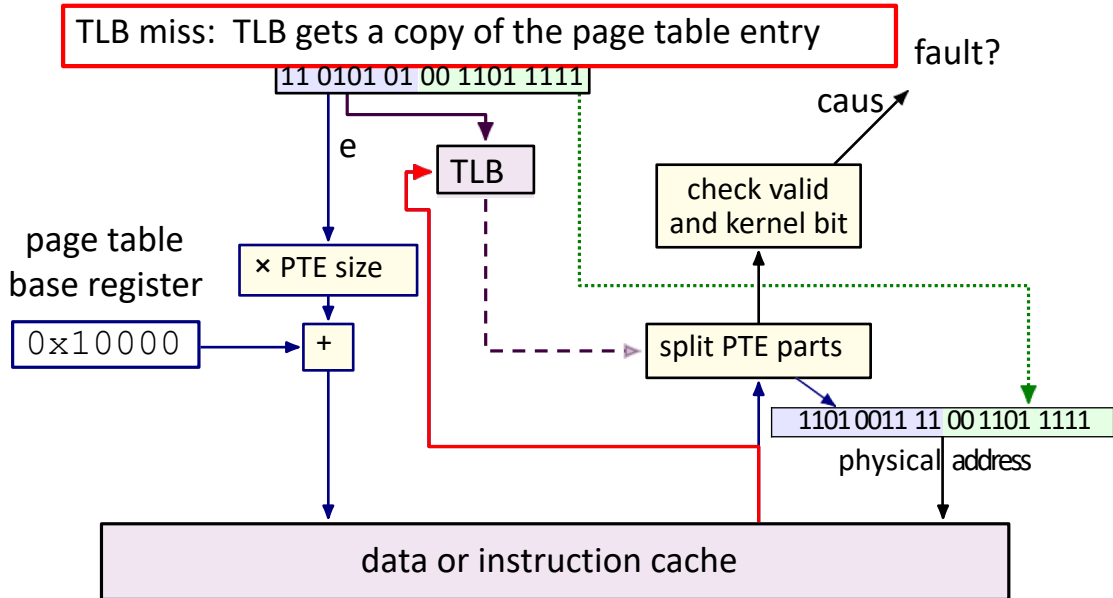
TLB and the MMU (2)



TLB and the MMU (2)



TLB and the MMU (2)



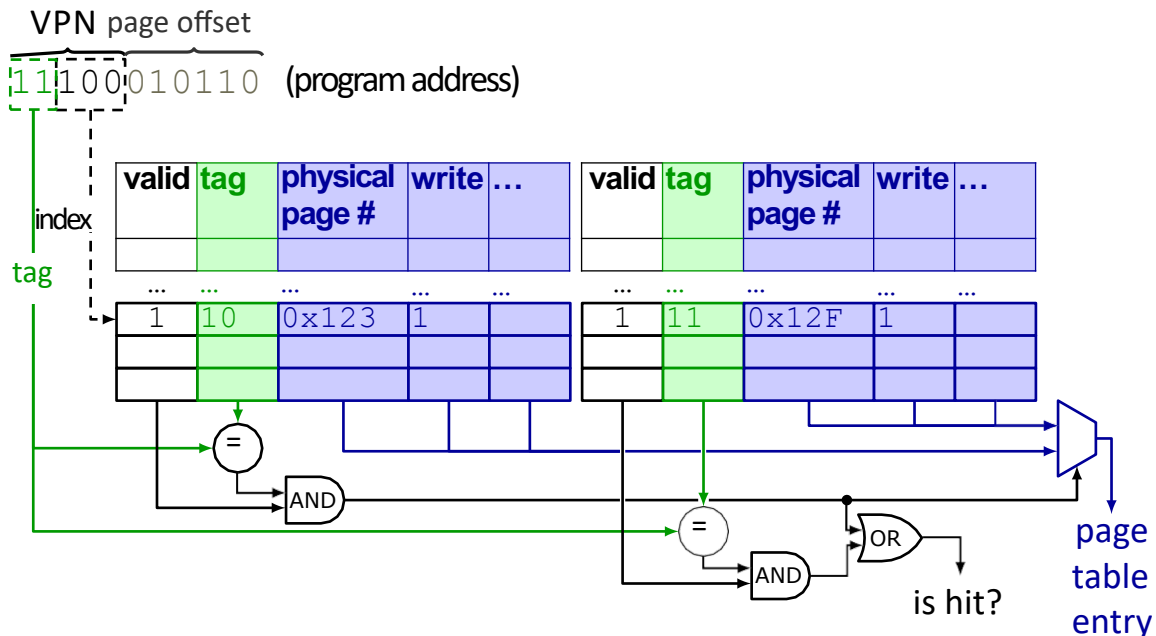
TLB and multi-level page tables

TLB caches **valid last-level page table entries**

doesn't matter which last-level page table

means TLB output can be used directly to form address

TLB organization (2-way set associative)



address splitting for TLBs

virtual address (48 bits)		
VPN (x bit)		page offset (y bit)
TLB tag (x1 bit)	TLB index (x2 bit)	page offset (y bit)

address splitting for TLBs (1)

my desktop:

4KB (2^{12} byte) pages; 48-bit virtual address

64-entry, 4-way L1 data TLB

TLB index bits?

TLB tag bits?

address splitting for TLBs (1)

my desktop:

4KB (2^{12} byte) pages; 48-bit virtual address 64-entry,

4-way L1 data TLB

TLB index bits?

$$\mathbf{64/4 = 16}$$
 sets — 4 bits

TLB tag bits?

$$\mathbf{48 - 12 = 36}$$
 bit virtual address — $\mathbf{36 - 4 = 32}$ bit TLB tag

address splitting for TLBs (2)

my desktop:

4KB (2^{12} byte) pages; 48-bit virtual address

1536-entry ($3 \cdot 2^9$), 12-way L2 TLB

TLB index bits?

TLB tag bits?

address splitting for TLBs (2)

my desktop:

4KB (2^{12} byte) pages; 48-bit virtual address 1536-

entry ($3 \cdot 2^9$), 12-way L2 TLB

TLB index bits?

$$\mathbf{1536 / 12 = 128}$$
 sets — 7 bits

TLB tag bits?

$$\mathbf{48 - 12 = 36}$$
 bit virtual address — $\mathbf{36 - 7 = 29}$ bit TLB tag

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address $0x12345678$

virtual address (32 bits)		
VPN part1 (x1 bits)	VPN part2 (x2 bits)	page offset (y bits)
TLB tag (z1 bits)	TLB index (z2 bits)	page offset (y bits)

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address $0x12345678$

virtual address (32 bits)		
VPN part1 (8 bits)	VPN part2 (11 bits)	page offset (13 bits)
TLB tag (12 bits)	TLB index (7 bits)	page offset (13 bits)

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address `0x12345678`

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address `0x12345678`

0001 0010 0011 0100 0101 0110 0111 1000

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address 0x12345678

0001 0010 0011 0100 0101 0110 0111 1000

13-bit page offset 1 0110 0111 1000

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address $0x12345678$

0001 0010 0011 0100 0101 0110 0111 1000

13-bit page offset 1 0110 0111 1000

32 - **13** = **19**-bit VPN 0001 0010 0011 0100

010

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address $0x12345678$

0001 0010 0011 0100 0101 0110 0111 1000

13-bit page offset 1 0110 0111 1000

32 - **13** = **19**-bit VPN 0001 0010 0011 0100

010

8-bit first part of VPN 0001 0010

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; **2048** in second

split the address $0x12345678$

0001 0010 **0011 0100 0101** 0110 0111 1000

13-bit page offset 1 0110 0111 1000

32 - 13 = 19-bit VPN 0001 0010 0011 0100

010

8-bit first part of VPN 0001 0010

11-bit second part of VPN **0011 0100 010**

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address $0x12345678$

0001 0010 0011 0100 0101 0110 0111 1000

13-bit page offset 1 0110 0111 1000

32 - 13 = 19-bit VPN 0001 0010 0011 0100

010

8-bit first part of VPN 0001 0010

11-bit second part of VPN 0011 0100 010

7-bit TLB index 0100 010

address splitting exercise (3)

384-entry, 3-way set-associative TLB

32-bit virtual address; 8KB pages 2-

level page table; 4 byte PTEs

256 entries in first level; 2048 in second

split the address $0x12345678$

0001 0010 0011 0100 0101 0110 0111 1000

13-bit page offset 1 0110 0111 1000

32 - 13 = 19-bit VPN 0001 0010 0011 0100

010

8-bit first part of VPN 0001 0010

11-bit second part of VPN 0011 0100 010

7-bit TLB index 0100 010

19 - 7 = 12-bit TLB tag 0001 0010 0011

TLB access pattern

64-byte pages

8 entries, 4

index	V	tag	PTE sets	V	tag	PTE	LRU
00							
01							
10							
11							

address (hex)	hit?
0001 00 000000 (100)	
1101 00 000001 (D01)	
0001 00 001010 (10A)	
1101 00 100001 (D21)	
0000 11 111100 (0FC)	
1100 11 111000 (CF8)	
1111 00 101000 (F23)	

VPN

changing page tables

what happens to TLB when page table base pointer is changed?

e.g. context switch

most entries in TLB refer to things from **wrong process**

oops — read from the wrong process's stack?

changing page tables

what happens to TLB when page table base pointer is changed?

e.g. context switch

most entries in TLB refer to things from **wrong process**

oops — read from the wrong process's stack?

option 1: **invalidate** all TLB entries

side effect on “change page table base register” instruction

changing page tables

what happens to TLB when page table base pointer is changed?

e.g. context switch

most entries in TLB refer to things from **wrong process**

oops — read from the wrong process's stack?

option 1: **invalidate** all TLB entries

side effect on “change page table base register” instruction

option 2: TLB entries contain process ID

set by OS (special register)

checked by TLB in addition to TLB tag, valid bit

editing page tables

what happens to TLB when OS changes a page table entry?

invalid to valid — nothing needed

TLB doesn't contain invalid entries

MMU will check memory again

valid to invalid — OS needs to tell processor to invalidate it

special instruction (x86: `invlpg`)


valid to other valid — OS needs to tell processor to invalidate it

TLB shutdown


program A pages



mark evicted page invalid in page table



program B pages

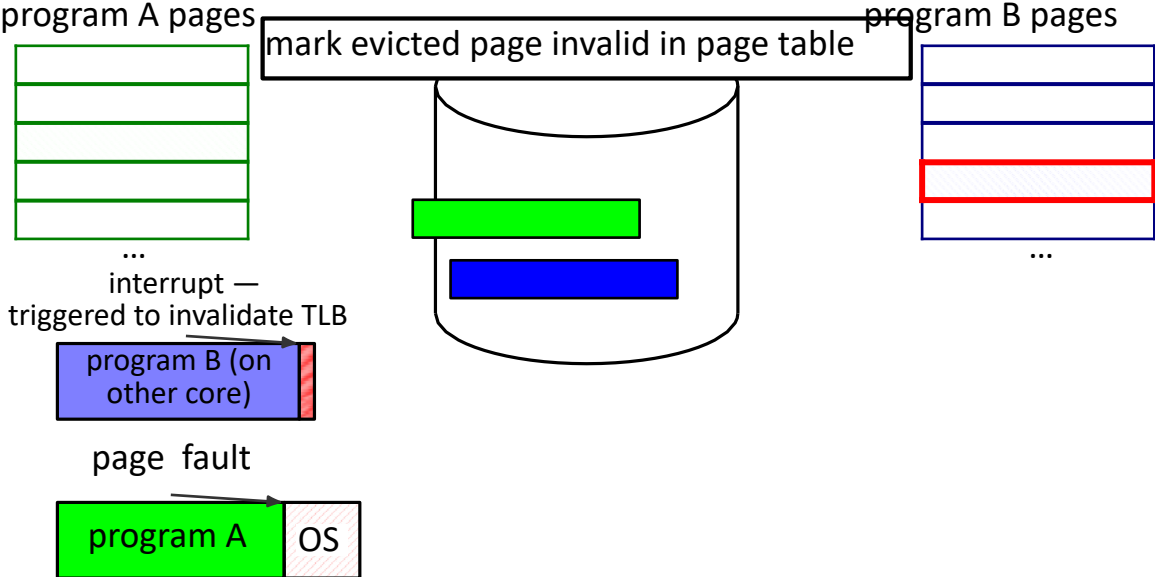


program B (on other core)

page fault




TLB shutdown



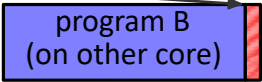
TLB shutdown

program A pages




...

interrupt —
triggered to invalidate TLB




page fault

start read



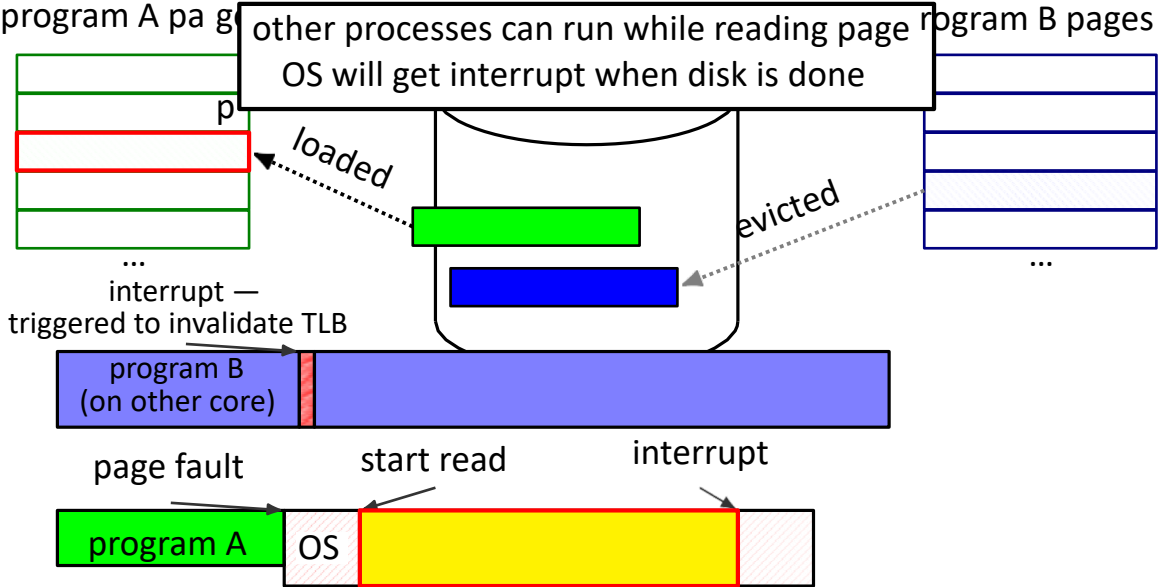
evicted

program B pages




...

TLB shutdown




TLB shutdown


program A pages



process A's page table updated and restarted from point of fault



program B pages



interrupt — triggered to invalidate TLB



page fault

start read

interrupt



loaded

evicted

x86-64 page table entries (1)

6	6	6	6	5	5	5	5	5	5	5	5	5	5	5		M ¹	M-1			3	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	0		
3	2	1	0	9	8	7	6	5	4	3	2	1								2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0				
X	Prot. Key ⁴	Ignored	Rsvd.	Address of 4KB page frame																Ign.	G	P	A	D	A	P	C	D	P	W	T	U	/	S	R	/	W	1	PTE: 4KB page							
Ignored																										0	PTE: not present																			

present = valid

R/W = writes allowed?

U/S = kernel-only? (“user/supervisor”)

XD = execute-disable?

A = accessed? (MMU sets to 1 on page read/write) D =

dirty? (MMU sets to 1 on page write)

helps support writeback policy for swapping

x86-64 page table entries (2)

666655555555		M ¹ M-1		333222222222211111111111																										
3210987654321				210987654321098765432109876543210																										
X D	Prot. Key ⁴	Ignored	Rsvd.	Address of 4KB page frame																Ign.	G	P A T	D	A	P C D	P W T	U / S	R / W	1	PTE: 4KB page
Ignored																			0	PTE: not present										

G = global? (shared between all page tables)

PWT, PCD, PAT = control how caches work when accessing physical page:
 can disable using the cache entirely
 can disable write-back (use write-through instead) multicore-related cache settings
 (and some other settings)

book's diagram

