


Page Tables

2025 Winter ECE353: Systems Software
Jon Eyolfson

Lecture 8
2.0.0

Processes Use A Register Like `satp` to Set the Root Page Table



Alignment: Memory Eventually Lines Up with Byte 0

If pages are *4096 byte aligned* in memory it means pages always start when the lower 12 bits are zero, in computing we like alignment

If a page started at address `0x7C00` its last byte would be at address `0x8BFF`

Instead, a page would start at `0x7000` and end at `0x7FFF`

Question: Is address `0xEC` 8 byte aligned?

Consider Translating a Two-Level Page Table Additional Level

Assume our process uses just one virtual address at `0x3FFFF008`
or `0b11_1111_1111_1111_1111_0000_0000_1000`
or `0b111111111_111111111_000000001000`

We'll just consider a 30-bit virtual address with a page size of 4096 bytes
We would need a 2 MiB page table if we only had one ($2^{18} \times 2^3$)

Instead, we have a 4 KiB L1 page table ($2^9 \times 2^3$) and a 4 KiB L0 page table
Total of 8 KiB instead of 2 MiB

Note: worst case if we used all virtual addresses we would consume 2 MiB +
4 KiB

Translating 3FFFF008 with a Two-Level Page Table

Consider the L1 table with the entry:

Index	PPN
511	0x8

Consider the L0 table located at 0x8000 with the entry:

Index	PPN
511	0xCAFE

The final translated physical address would be: 0xCAFE008

Let's Simulate an MMU

lectures/07-virtual-memory in the materials repository

Remember each process would have its own unique root page table

How Many Page Tables Do We Need?

Let's assume our program uses 512 pages

What's the minimum number of page tables we need?

What's the maximum number of page tables?

How Many Levels Do I Need?

Assume we have a 32-bit virtual address with a page size of 4096 bytes and a PTE size of 4 bytes

We want each page table to fit into a single page

Find the number of PTEs we could have in a page (2^{10})

$\log_2(\text{\#PTEs per Page})$ is the number of bits to index a page table

$$\text{\#Levels} = \lceil \frac{\text{Virtual Bits} - \text{Offset Bits}}{\text{Index Bits}} \rceil$$

How Many Levels Do I Need?

Assume we have a 32-bit virtual address with a page size of 4096 bytes and a PTE size of 4 bytes

We want each page table to fit into a single page

Find the number of PTEs we could have in a page (2^{10})

$\log_2(\text{\#PTEs per Page})$ is the number of bits to index a page table

$$\text{\#Levels} = \lceil \frac{\text{Virtual Bits} - \text{Offset Bits}}{\text{Index Bits}} \rceil$$

$$\text{\#Levels} = \lceil \frac{32 - 12}{10} \rceil = 2$$

Using the Page Tables for Every Memory Access is Slow

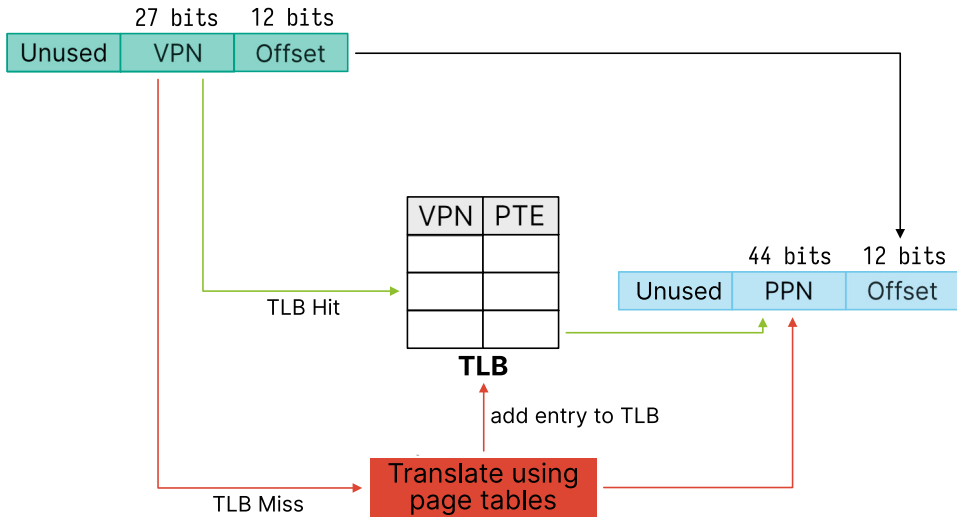
We need to follow pointers across multiple levels of page tables!

We'll likely access the same page multiple times
(close to the first access time)

A process may only need a few VPN \rightarrow PPN mappings at a time

Our solution is another computer science classic: caching

A Translation Look-Aside Buffer (TLB) Caches PTEs



Effective Access Time (EAT)

Assume a single page table

(there's only one additional memory access in the page table)

$$\text{TLB_Hit_Time} = \text{TLB_Search} + \text{Mem}$$

$$\text{TLB_Miss_Time} = \text{TLB_Search} + 2 \times \text{Mem}$$

$$\text{EAT} = \alpha \times \text{TLB_Hit_Time} + (1 - \alpha) \times \text{TLB_Miss_Time}$$

If $\alpha = 0.8$, $\text{TLB_Search} = 10$ ns, and accesses take 100 ns, calculate EAT

$$\text{EAT} = 0.8 \times 110 \text{ ns} + 0.2 \times 210 \text{ ns}$$

$$\text{EAT} = 130 \text{ ns}$$

Context Switches Require Handling the TLB

You can either flush the cache, or attach a process ID to the TLB

Most implementation just flush the TLB

RISC-V uses a `sfence.vma` instruction to flush the TLB

On x86 loading the base page table will also flush the TLB

TLB Testing

Check out `lectures/08-page-tables/test-tlb`

(you may need to `git submodule update --init --recursive`)

```
./test-tlb <size> <stride>
```

Creates a `<size>` memory allocation and accesses it every `<stride>` bytes

Results from my laptop:

```
> ./test-tlb 4096 4
  1.93ns (~7.5 cycles)
> ./test-tlb 536870912 4096
155.51ns (~606.5 cycles)
> ./test-tlb 16777216 128
 14.78ns (~57.6 cycles)
```

Use `sbrk` for Userspace Allocation

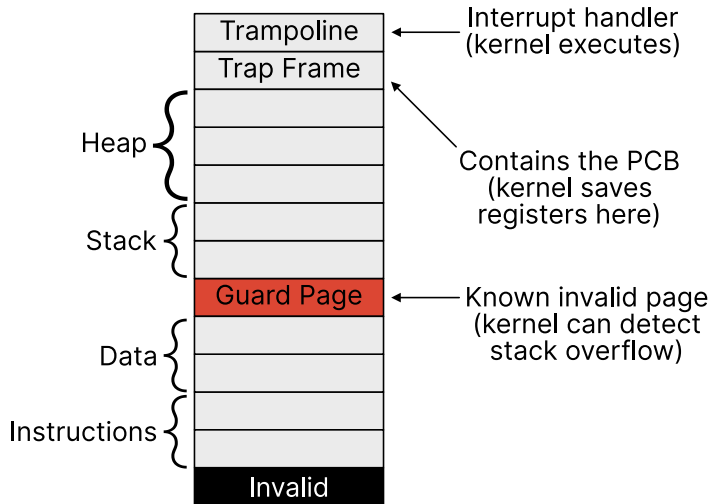
This call grows or shrinks your heap (the stack has a set limit)

For growing, it'll grab pages from the free list to fulfill the request
The kernel sets `PTE_V` (valid) and other permissions

In memory allocators this is difficult to use, you'll rarely shrink the heap
It'll stay claimed by the process, and the kernel cannot free pages

Memory allocators use `mmap` to bring in large blocks of virtual memory

The Kernel Initializes the Process' Address Space



The Kernel Can Provide Fixed Virtual Addresses

It allows the process to access kernel data without using a system call

For instance `clock_gettime` does not do a system call

It just reads from a virtual address mapped by the kernel

Page Faults Allow the Operating System to Handle Virtual Memory

Page faults are a type of exception for virtual memory access

Generated if it cannot find a translation, or permission check fails

This allows the operating system to handle it

We could lazily allocate pages, implement copy-on-write, or swap to disk

Page Tables Translate Virtual to Physical Addresses

The MMU is the hardware that uses page tables, which may:

- Be a single large table (wasteful, even for 32-bit machines)
- Use the kernel allocated pages from a free list
- Be a multi-level to save space for sparse allocations
- Use a TLB to speed up memory accesses