

ITP 30002 Operating System

Paging: Smaller Tables

OSTEP Chapter 20

Shin Hong

Memory-efficient Page Table Structures

2

- Array-based page tables take up too much memory even if most page entries are invalid
 - e.g., for a 32-bit address space with 4 KB pages, a per-process page table takes 4 MB
- Approaches
 1. use bigger pages
 2. per-segment page tables
 3. multi-level page tables
 4. inverted page table


Paging: Smaller
Page Tables

ITP 30002
Operating System

2023-04-13

Per-segment Page Table

3



PFN	valid	prot	present	dirty
10	1	r-x	1	0
-	0	—	-	-
-	0	—	-	-
-	0	—	-	-
23	1	rw-	1	1
-	0	—	-	-
-	0	—	-	-
-	0	—	-	-
-	0	—	-	-
-	0	—	-	-
-	0	—	-	-
-	0	—	-	-
-	0	—	-	-
-	0	—	-	-
28	1	rw-	1	1
4	1	rw-	1	1

- Motivation: it is likely that only first few pages of each segment is used
- Allocate a variable-length a page table for each segment
 - re-use base-bound register pairs
 - a base register points to the beginning of a page table, and a bound register represents the number of allocated pages in the segment


Paging: Smaller
Page Tables

ITP 30002
Operating System

2023-04-13

Example

4



- use the two bits to represent a segment identifier
 - 01 for code, 10 for the heap, 11 for the stack
- access a segment page table at a TLB miss

```
SN          = (VirtualAddress & SEG_MASK) >> SN_SHIFT
VPN         = (VirtualAddress & VPN_MASK) >> VPN_SHIFT
AddressOfPTE = Base[SN] + (VPN * sizeof(PTE))
```

		PFN	valid	prot	present	dirty
		10	1	r-x	1	0
Base[01]	→					
		PFN	valid	prot	present	dirty
		23	1	rw-	1	1
Base[10]	→					
		PFN	valid	prot	present	dirty
		28	1	rw-	1	1
Base[11]	→					
		4	1	rw-	1	1

Paging: Smaller
Page Tables

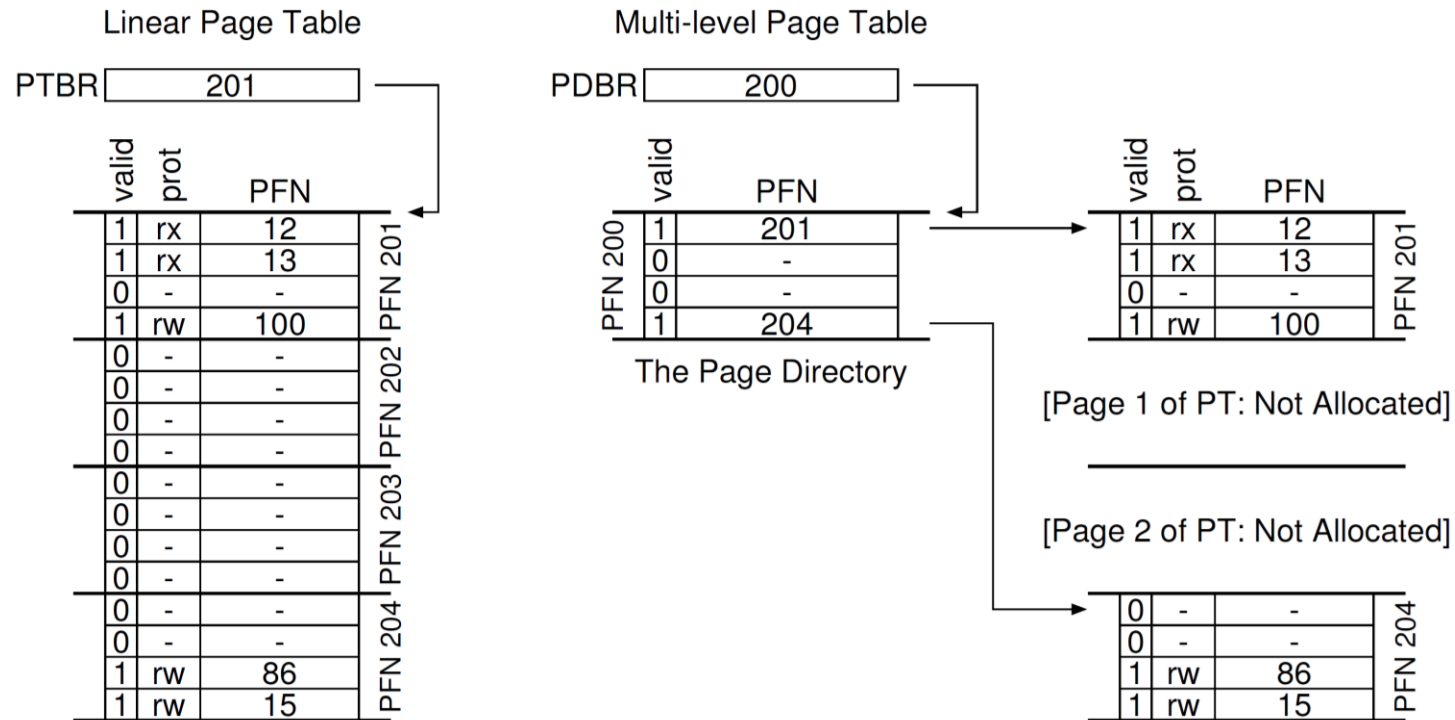
ITP 30002
Operating System

2023-04-13

Multi-level Page Tables

5

- Divide a page table into page-size units
 - A page table spans over multiple pages
 - Do not allocate a page if the corresponding piece of a page table has no valid entry
- Create a page directory as an index of the allocated pages for a page table
 - E.g.




Paging: Smaller
Page Tables

ITP 30002
Operating System

2023-04-13

6

- | | |
|-----------|------------------|
| 0000 0000 | code |
| 0000 0001 | code |
| 0000 0010 | (free) |
| 0000 0011 | (free) |
| 0000 0100 | heap |
| 0000 0101 | heap |
| 0000 0110 | (free) |
| 0000 0111 | (free) |
| | ... all free ... |
| 1111 1100 | (free) |
| 1111 1101 | (free) |
| 1111 1110 | stack |
| 1111 1111 | stack |



2023-04-13

Two-level Page Table Control Flow

7

```
1  VPN = (VirtualAddress & VPN_MASK) >> SHIFT
2  (Success, TlbEntry) = TLB_Lookup(VPN)
3  if (Success == True)    // TLB Hit
4      if (CanAccess(TlbEntry.ProtectBits) == True)
5          Offset = VirtualAddress & OFFSET_MASK
6          PhysAddr = (TlbEntry.PFN << SHIFT) | Offset
7          Register = AccessMemory (PhysAddr)
8      else
9          RaiseException(PROTECTION_FAULT)
10 else    // TLB Miss
11     // first, get page directory entry
12     PDIndex = (VPN & PD_MASK) >> PD_SHIFT
13     PDEAddr = PDBR + (PDIndex * sizeof(PDE))
14     PDE = AccessMemory (PDEAddr)
15     if (PDE.Valid == False)
16         RaiseException(SEGMENTATION_FAULT)
17     else
18         // PDE is valid: now fetch PTE from page table
19         PTIndex = (VPN & PT_MASK) >> PT_SHIFT
20         PTEAddr = (PDE.PFN << SHIFT) + (PTIndex * sizeof(PTE))
21         PTE = AccessMemory (PTEAddr)
22         if (PTE.Valid == False)
23             RaiseException(SEGMENTATION_FAULT)
24         else if (CanAccess(PTE.ProtectBits) == False)
25             RaiseException(PROTECTION_FAULT)
26         else
27             TLB_Insert(VPN, PTE.PFN, PTE.ProtectBits)
28             RetryInstruction()
```

Paging: Smaller
Page Tables

ITP 30002
Operating System

2023-04-13

Inverted Page Table

8

- Keep a single page table that has an entry for each frame
 - each frame is mapped to a pair of a process ID and a VPN
 - one page table is enough for a single system
- Searching the entry for a VPN of a process consumes much more time than array-based page tables
 - linear search, hashing, etc.

Paging: Smaller
Page Tables


ITP 30002
Operating System

2023-04-13

Intel x86-64

9

- 64 bits is ginormous (> 16 exabytes)
- In practice only implement 48 bit addressing
 - Page sizes of 4 KB, 2 MB, 1 GB
 - Four levels of paging hierarchy
- Can also use PAE so virtual addresses are 48 bits and physical addresses are 52 bits



Paging: Smaller
Page Tables

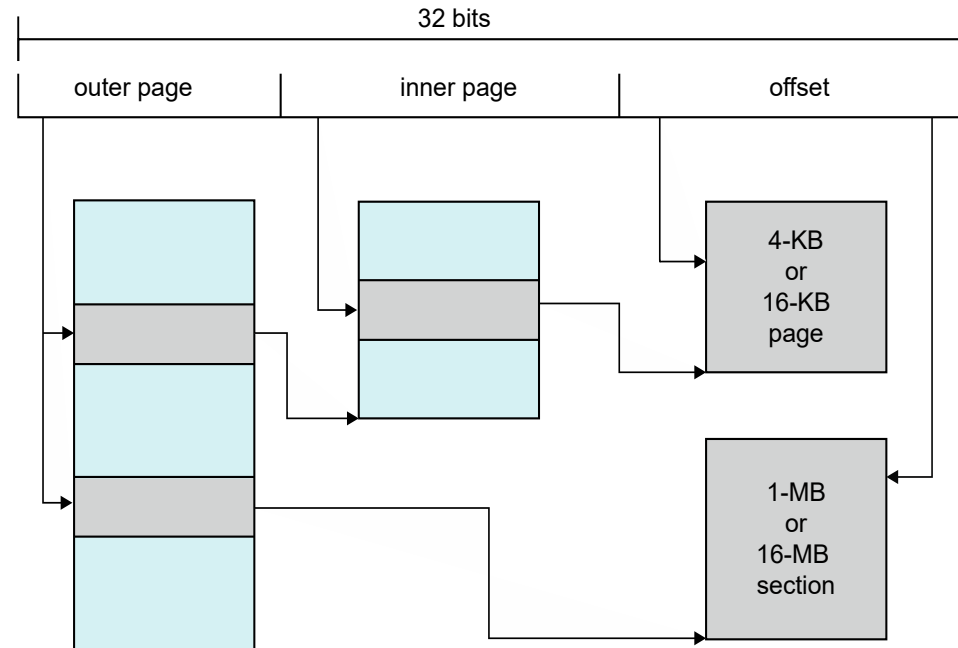
ITP 30002
Operating System

2023-04-13

ARM Architecture

10

- Modern, energy efficient, 32-bit CPU for mobile platform
- 4 KB and 16 KB pages
- 1 MB and 16 MB sections (large-size pages)
- two-level paging for pages & One-level paging for sections
- Two levels of TLBs
 - Inner is single main TLB
 - Outer level has two micro TLBs (one data, one instruction)
 - First inner is checked, on miss others are checked, and on miss page table walk performed by CPU



Paging: Smaller
Page Tables

ITP 30002
Operating System

2023-04-13