

計算機アーキテクチャ 第一 (E)



10. 主記憶とファイルメモリの管理, 多重仮想記憶, 記憶保護

吉瀬 謙二 計算工学専攻
kise_at_cs.titech.ac.jp
W641講義室 木曜日13:20 – 14:50

Acknowledgement

- Lecture slides for Computer Organization and Design, Third Edition, courtesy of **Professor Mary Jane Irwin**, Penn State University
- Lecture slides for Computer Organization and Design, third edition, Chapters 1-9, courtesy of **Professor Tod Amon**, Southern Utah University.

Adapted from *Computer Organization and Design*, Patterson & Hennessy, © 2005



Reliability(信頼性), Availability


- Reliability – measured by the **mean time to failure** (平均故障寿命, MTTF). Service interruption is measured by **mean time to repair** (平均修復時間, MTTR)
- Availability(アベイラビリティ)**

$$\text{Availability} = \text{MTTF} / (\text{MTTF} + \text{MTTR})$$
- To increase MTTF, either improve the quality of the components or design the system to continue operating in the presence of faulty components
 - Fault avoidance: preventing fault occurrence by construction
 - Fault tolerance: using redundancy to correct or bypass faulty components (hardware)

7

RAID: Disk Arrays

Redundant Array of Inexpensive Disks



- Arrays of small and inexpensive disks
 - Increase potential **throughput** by having many disk drives
 - Data is spread over multiple disk
 - Multiple accesses are made to several disks at a time
- Reliability** is lower than a single disk
- But **availability** can be improved by adding **redundant disks (RAID)**

8


RAID: Level 0 (冗長性なし; Striping ストライピング)



- Multiple smaller disks as opposed to one big disk
 - Spreading the blocks over multiple disks – **striping** – means that multiple blocks can be accessed in parallel increasing the performance
 - A 4 disk system gives four times the throughput of a 1 disk system
 - Same cost as one **big disk** – assuming 4 small disks cost the same as one big disk
- No redundancy, so what if one disk fails?

9

RAID: Level 1 (Redundancy via Mirroring)



- Uses twice as many disks for redundancy so there are always two copies of the data
 - The number of redundant disks = the number of data disks so twice the cost of one big disk
 - writes have to be made to both sets of disks, so writes would be only 1/2 the performance of RAID 0
- What if one disk fails?
 - If a disk fails, the system just goes to the "mirror" for the data

10

RAID: Level 0+1 (Striping with Mirroring)



- Combines the best of RAID 0 and RAID 1, data is striped across four disks and mirrored to four disks
 - Four times the throughput (due to striping)
 - # redundant disks = # of data disks so twice the cost of one big disk
 - writes have to be made to both sets of disks, so writes would be only 1/2 the performance of RAID 0
- What if one disk fails?
 - If a disk fails, the system just goes to the "mirror" for the data







11

RAID: Level 4 (Block-Interleaved Parity)




- Cost of higher availability still only 1/N but the parity is stored as **blocks** associated with sets of data blocks
 - Four times the throughput (striping)
 - # redundant disks = 1 × # of protection groups
 - Supports "small reads" and "small writes" (reads and writes that go to just one (or a few) data disk in a protection group)

12



Virtual Memory (仮想記憶)

- Use main memory as a “cache” for secondary memory
 - Simplifies loading a program for execution by providing for code relocation (i.e., the code can be loaded anywhere in main memory)
 - Provides the ability to easily run programs larger than the size of physical memory
 - Allows efficient and safe sharing of memory among multiple programs



19


Virtual Memory (仮想記憶)

- What makes it work? – again the **Principle of Locality**
 - A program is likely to access a relatively small portion of its address space during any period of time

20

Virtual Memory (仮想記憶)


- Each program is compiled into its own address space – a “virtual” address space
 - During run-time each **virtual address, VA** (仮想アドレス) must be translated to a **physical address, PA** (物理アドレス)



21

Two Programs Sharing Physical Memory

- A program’s address space is divided into **pages** (all one fixed size) or **segments** (variable sizes)
 - The starting location of each page (either in **main memory** or in **secondary memory**) is contained in the program’s **page table**




22

Address Translation

- A virtual address is translated to a physical address by a combination of hardware and software


Virtual Address (VA)





- So each memory request first requires an **address translation** from the virtual space to the physical space

23


Address Translation Mechanisms




24




26




27



28




29



30

A TLB in the Memory Hierarchy



- **page fault** : page is not in physical memory
- **TLB misses** are much more frequent than true page faults

31

Two Machines' TLB Parameters

	Intel P4	AMD Opteron
TLB organization	1 TLB for instructions and 1 TLB for data Both 4-way set associative Both use ~LRU replacement	2 TLBs for instructions and 2 TLBs for data Both L1 TLBs fully associative with ~LRU replacement Both L2 TLBs are 4-way set associative with round-robin LRU
Both have 128 entries	Both L1 TLBs have 40 entries Both L2 TLBs have 512 entries	TLB misses handled in hardware
TLB misses handled in hardware	TLB misses handled in hardware	TLB misses handled in hardware

32


TLB Event Combinations

TLB	Page Table	Cache	Possible? Under what circumstances?
Hit	Hit	Hit	Yes – what we want!
Hit	Hit	Miss	Yes – although the page table is not checked if the TLB hits
Miss	Hit	Hit	Yes – TLB miss, PA in page table
Miss	Hit	Miss	Yes – TLB miss, PA in page table, but data not in cache
Miss	Miss	Miss	Yes – page fault
Hit	Miss	Miss/Hit	Impossible – TLB translation not possible if page is not present in memory
Miss	Miss	Hit	Impossible – data not allowed in cache if page is not in memory

33


Reducing Translation Time

- Can **overlap** the cache access with the TLB access
 - Works when the high order bits of the VA are used to access the TLB while the low order bits are used as index into cache



34

A TLB in the Memory Hierarchy




- **page fault** : page is not in physical memory
- **TLB misses** are much more frequent than true page faults

35

Why Not a **Virtually Addressed Cache**?

- A **virtually addressed cache** would only require address translation on cache misses



but

- Two different virtual addresses can map to the same physical address (when processes are sharing data),
- Two different cache entries hold data for the same physical address – **synonyms (別名)**
 - Must update all cache entries with the same physical address or the memory becomes inconsistent

36



The Hardware/Software Boundary

- What parts of the virtual to physical address translation is done by or assisted by the hardware?
 - **Translation Lookaside Buffer (TLB)** that caches the recent translations
 - TLB access time is part of the cache hit time
 - May cause an extra stage in the pipeline for TLB access
 - Page table storage, fault detection and updating
 - **Page faults** result in **interrupts (precise)** that are then handled by the **OS**
 - Hardware must support (i.e., update appropriately) **Dirty** and **Reference bits** (e.g., **-LRU**) in the Page Tables

37



Summary

- The Principle of Locality:
 - Program likely to access a relatively small portion of the address space at any instant of time.
 - **Temporal Locality:** Locality in Time
 - **Spatial Locality:** Locality in Space
- Caches, TLBs, Virtual Memory all understood by examining how they deal with the four questions
 1. Where can block be placed?
 2. How is block found?
 3. What block is replaced on miss?
 4. How are writes handled?
- **Page tables** map virtual address to physical address
 - **TLBs** are important for fast translation

38