CS152/252A
Spring 2026

Memory Hierarchy
Discussion 4

1 Prefetching Basics

1.1 Define and give a formula to compute each prefetching metric.

Accuracy:

Coverage:

Timeliness:

1.2 What are the memory access patterns for:

Instruction prefetching:

Data prefetching:

1.3 Which prefetching algorithm(s) are most appropriate for:

Instruction memory:

Data memory:

2 Software Prefetching

For the following program, assume 128B cache lines (each row fits entirely in a cache line). Without the prefetch, the inner loop takes 50 cycles. The L1 miss penalty is 40 cycles. What should OFFSET be to minimize the total program cycles?

```
int A[N][M]; // N=32, M=32
int sum = 0;

for (int j = 0; j < M; j++) {
  for (int i = 0; i < N; i++) {
    // prefetches from (A + M*i + j + OFFSET)
    prefetch(&A[i][j] + OFFSET);
    sum += A[i][j];
  }
}
```

3 Linear vs Hierarchical Page Tables

3.1 Consider 4 GiB (32-bit) of addressable virtual memory, 4 KiB pages, 4-byte PTEs.

Bits in the page offset:

Bits in the virtual page number:

Number of pages:

3.2 Consider a linear page table for a process with only 1 page mapped to physical memory (paged in)

Number of valid page table entries (PTEs):

Total size of the page table:

3.3 Consider a 2-level page table for a process with only 1 page mapped to physical memory (paged in). Assume that VPN bits are split equally between the two levels.

Number of valid page table entries (PTEs):

Total size of the page table structure:

4 Page Size Tradeoffs

What are the benefits of a larger page size? What are the benefits of a smaller page size?

5 Hierarchical Page Table Example

Assume: 8-bit virtual addresses, 32-bit words, 32-bit PPNs, 16-byte pages, two-level page table, LRU 4-entry TLB. At the beginning, the TLB is empty and the free pages list contains 0x9, 0x5, 0xA, 0x7, 0x1, 0x3, 0xB, 0xD, 0xE, and 0xF in that order. PTBR is set to 0.

5.1 How many bytes of virtual memory are addressable?

5.2 How many bytes of physical memory are addressable?

5.3 Why might DRAM size > virtual address space size be useful?

5.4 Divide the virtual address into:

Offset bits:

Second level index bits:

Top level index bits:

5.5 How many entries are in the…

Top level table?

Second level tables (each)?

5.6 Fill in the tables. The first address translation is given as an example.

Free pages: 0x9, 0x5, 0xA, 0x7, 0x1, 0x3, 0xB, 0xD, 0xE, 0xF

Execution Behavior:

Virtual Address	Index1	Index2	TLB hit / miss	Page hit / page fault	Physical address
0x68	0x1	0x2	miss	hit	0x128
0x14					
0x6C					
0x90					
0x74					
0xE4					
0x18					
0xD0					

TLB:

VPN	0x6			
PPN	0x12			

Memory Contents:

Address	Content
0x00	0x06
0x04	0x04
0x08	0x02
0x0C	
0x10	
0x14	
0x18	
0x1C	
0x20	0x08
0x24	
0x28	
0x2C	
0x30	
0x34	
0x38	
0x3C	
0x40	
0x44	
0x48	0x12
0x4C	0x11
0x50	
0x54	
0x58	
0x5C	
0x60	
0x64	0x13
0x68	
0x6C	