


ECE 361 Homework 1
Fall 2004
Due: 10/14/04

1. What is the approximate cost of a die in the wafer shown in Figure 1? Assume that an 8-inch costs \$1000 and that the defect density is 1 per square centimeter. Use the number of dies per wafer given in the figure caption.

Some Necessary Equations:

$$\text{Cost per die} = \frac{\text{Cost per wafer}}{\text{Dies per wafer} * \text{yield}}$$

$$\text{Yield} = \frac{1}{\left(1 + \frac{\text{Defects per area} * \text{Die area}}{2}\right)^2}$$



(Fig 1) An 8 inch (200-mm) diameter wafer containing Intel Pentium processors. The number of Pentium dies per wafer at 100% yield is 196. The die area is 91 mm^2 , and it contains about 3.3 million transistors.

2. DRAM chips have significantly increased in die size with each generation, yet yields have stayed about the same (43% to 48%). Figure 2 shows key statistics for DRAM production over the years.

Given the increase in die area of DRAMs, what parameter (see the equations) must improve to maintain yield?

(Fig 2)

Year	Capacity (Kbits)	Die area (sq. cm)	Wafer diameter (inches)	Yield
1980	64	0.16	5	48%
1983	256	0.24	5	46%
1985	1024	0.42	6	45%
1989	4096	0.65	6	43%
1992	16384	0.97	8	48%

3. Consider two different implementations, M1 and M2, of the same instruction set. There are four classes of instructions (A, B, C, and D) in the instruction set.

M1 has a clock rate of 500 MHz. The average number of cycles for each instruction class on M1 is as follows:

Class	CPI for this class
A	1
B	2
C	3
D	4

M2 has a clock rate of 750 MHz. The average number of cycles for each instruction class on M2 is as follows:

Class	CPI for this class
A	2
B	2
C	4
D	4

Assume that peak performance is defined as the fastest rate that a machine can execute an instruction sequence chosen to maximize that rate. What are the peak performances of M1 and M2 expressed as instructions per second?

- If the number of instructions executed in a certain program is divided equally among the classes of instructions in Problem 3, how much faster is M2 than M1?
- Assuming the CPI values from Problem 3 and the instruction distribution from Problem 4, at what clock rate would M1 have the same performance as the 750-MHz version of M2?
- The table below shows the number of floating-point operations executed in two different programs and the runtime for those programs on three different machines:

Program	Floating-point operations	Execution time in Seconds		
		Computer A	Computer B	Computer C
Program 1	10,000,000	1	10	20
Program 2	100,000,000	1000	100	20

Which machine is fastest according to total execution time? How much faster is it than the other two machines?

- Suppose we have made the following measurements of average CPI for instructions:

Instruction	Average CPI
Arithmetic	1.0 clock cycles
Data transfer	1.4 clock cycles
Conditional branch	1.7 clock cycles
Jump	1.2 clock cycles

Compute the effective CPI for MIPS. Average the instruction frequencies for gcc and spice in Figure 3 to obtain the instruction mix.

(Figure 3) MIPS instruction classes, examples, correspondence to high-level program language constructs, and percentage of MIPS instructions executed by category for two programs, gcc and spice.

Instruction class	MIPS examples	HLL correspondence	Frequency	
			gcc	Spice
Arithmetic	add, sub, addi	Operations in assignment statements	48%	50%
Data transfer	lw, sw, lb, sb, lui	References to data structure, such as arrays	33%	41%
Conditional	beq, bne,	if statements	17%	8%

branch	slt, slti	and loops		
Jump	j, jr, jal	Procedure calls, returns, and case/switch statements	2%	1%

(Figure 4) shows the percentage of the individual MIPS instructions executed.

Core MIPS	Name	Gcc (%)	Spice (%)
Add	Add	0	0
Add immediate	Addi	0	0
Add unsigned	Addu	9	10
Add immediate unsigned	Addiu	17	1
Subtract unsigned	Subu	0	1
And	And	1	0
And immediate	Andi	2	1
Shift left logical	Sll	5	5
Shift right logical	Srl	0	1
Load upper immediate	Lui	2	6
Load word	Lw	21	7
Store word	Sw	12	2
Load byte	Lb	1	0
Store byte	Sb	1	0
Branch on equal (zero)	Beq	9	3
Branch on not equal (zero)	Bne	8	2
Jump and link	Jal	1	1
Jump register	Jr	1	1
Set less than	Slt	2	0
Set less than immediate	Slti	1	0
Set less than unsigned	Sltu	1	0
Set less than immediate unsigned	Sltiu	1	0
FP add double	Add.d	0	4
FP subtract double	Sub.d	0	3
FP multiply double	Sul.d	0	5
FP divide double	div.d	0	2
Load word to FP single	l.s	0	24
Store word to FP single	s.s	0	9
Branch on FP true	Bc1t	0	1
Branch on FP false	Bc1f	0	1
FP compare double	c.x.d	0	1
Move to FP	Mtcl	0	2
Move from FP	Mfc2	0	2
Convert float integer	Cut	0	1

Shift right arithmetic	Sra	2	0
Load half	Lh	1	0
Branch less than zero	Bltz	1	0
Branch greater than zero	Bgez	1	0
Branch less or equal zero	Blez	0	1

7. In this exercise, we'll examine quantitatively the pros and cons of adding an addressing mode to MIPS that allows arithmetic instructions to directly access memory, as is found on the 80x86. The primary benefit is that fewer instructions will be executed because we won't have to first load a register. The primary disadvantage is that the cycle time will have to increase to account for the additional time to read memory. Consider adding a new instruction:

Addm \$t2, 100(\$t3) # \$t2 = \$t2 + Memory[\$t3+100]

Assume that the new instruction will cause the cycle time to increase by 10%. Use the instruction frequencies for the gcc benchmark from Figure 3, and assume that two-thirds of the data transfers are loads and the rest are stores. Assume that the new instruction affects only the clock speed, not the CPI. What percentage of loads must be eliminated for the machine with the new instruction to have at least the same performance?