

EECS 151/251A Homework 1

Due Monday, Feb 1th, 2021

Problem 1: Pareto Optimal Frontier

John did a design space exploration for his design of a digital widget and came up with the following table of results for maximum frequency in GHz, energy efficiency in nanoJoules per operation, and cost as chip area in mm². List those rows that represent design points that lie on the Pareto optimal frontier.

f_{max}	Energy	Cost
2.0	20	2.0
1.5	10	1.5
1.5	17	1.5
1.5	20	1.0
1.0	10	1.5
1.0	20	1.0
1.0	10	1.5
0.2	5	1.0

Solution:

Rows on the optimal frontier:

f_{max}	Energy	Cost
2.0	20	2.0
1.5	10	1.5
1.5	20	1.0
0.2	5	1.0

Problem 2: Moore's Law Implications

Let's imagine an alternate history for a moment. Back in 1908, the Ford Motor Company introduced the Model T to the world as the first affordable automotive. This pioneering vehicle sported a 4-cylinder engine outputting 20 horsepower and came in any color you wanted as long as it was black! While working on developing the next generation of automobiles in 1910, one of the engineers at the company noticed that if they shrank the pistons in the engines of the cars, they could get the same performance out of the engine while using a leaner fuel-air mixture. This would mean each engine block could now fit even more pistons, and therefore output more power. Upon hearing of this, Henry Ford made the prediction that every 2 years the number of pistons in automobile engines would double, a notion that became known as "Ford's Law". Following this trend, how many pistons

would today's gas-powered cars have if this were a real phenomenon in the automotive industry? Is this a realistic number of pistons to have in an engine? Barring the issues with sizing the piston (assume this world could make infinitesimally small pistons), what are some additional issues with having this many pistons in one engine?

Solution:

If this trend continued from 1910 to today, we would have seen **55 doublings** of engine cylinder counts. From the original 4 cylinders the Model T began with, today's cars would have 1.44×10^{17} cylinders. This is unrealistically large. Some potential issues are **the huge energy density in the engine block** and **the problems of transferring power from the cylinders to the transmission**.

Problem 3: Technology Survey

Take a look at your laptop/desktop computer. What processor is it using? Which technology node was it designed in (14 nm, 12 nm, 7 nm, etc.)? How many cores does it have? Does it have dedicated hardware accelerator blocks?

Solution:

Answers will vary. I personally have an i7-6700K processor which was made in 14 nm. It has 4 cores (8 threads) with a base core clock of 4 GHz. It has some Intel integrated graphics accelerators built in, but these are generally unused in favor of a dedicated graphics card.

Problem 4: Die Cost

You want to fabricate 300 mm wafers with $\alpha = 3$ and a defect per unit area of $0.005 / \text{mm}^2$.

- The die area is 25 mm^2 and the wafer cost is \$15k. What is your die yield and die cost?
- Your company is now pivoting to making FPGAs that are 500 mm^2 each, and the wafer costs have gone down so they are now \$8k each. What is your die yield and die cost?

Solution:

(a)

$$\text{Die Yield} = \left(1 + \frac{0.005 / \text{mm}^2 \cdot 25 \text{ mm}^2}{3}\right)^{-3} = 0.88$$

$$\text{Dies per wafer} = \frac{\pi \cdot (300 \text{ mm}/2)^2}{25 \text{ mm}^2} - \frac{\pi \cdot 300 \text{ mm}}{\sqrt{2} \cdot 25 \text{ mm}^2} = 2694$$

$$\text{Die Cost} = \frac{\$15,000}{2694 \cdot 0.88} = \$6.33$$

(b)


$$\text{Die Yield} = \left(1 + \frac{0.005 / \text{mm}^2 \cdot 500 \text{ cm}^2}{3}\right)^{-3} = 0.16$$

$$\text{Dies per wafer} = \frac{\pi \cdot (300 \text{ mm}/2)^2}{500 \text{ mm}^2} - \frac{\pi \cdot 300 \text{ mm}}{\sqrt{2} \cdot 500 \text{ mm}^2} = 112$$

$$\text{Die Cost} = \frac{\$8,000}{112 \cdot 0.16} = \$446.43$$

Problem 5. Boolean Logic

For the digital logic circuit shown below, give the truth table.



Solution:

Truth Table

A	B	C	D	Z
0	0	0	0	0
0	0	0	1	0
0	0	1	0	0
0	0	1	1	0
0	1	0	0	1
0	1	0	1	1
0	1	1	0	0
0	1	1	1	0
1	0	0	0	1
1	0	0	1	1
1	0	1	0	1
1	0	1	1	1
1	1	0	0	1
1	1	0	1	1
1	1	1	0	0
1	1	1	1	0

Notice that the final OR gate and the D input are redundant, so if this was given to a synthesizer, these would likely be optimized out.

Problem 6. Boolean Functions of 2 inputs

How many unique Boolean functions of 2 inputs exist? Give names for as many of them as you can, and write a Boolean expression for those you can't.

Solution:

There are 16 functions of 2 inputs. The common ones are ZERO, ONE, AND, NAND, OR, NOR, XOR, XNOR. The remaining functions are IS A, IS B, NOT A, NOT B, ONLY A, ONLY B, NOT ONLY A, NOT ONLY B (don't necessarily have to use these names of course).

- IS A: $Z = A$
- IS B: $Z = B$
- NOT A: $Z = \bar{A}$
- NOT B: $Z = \bar{B}$
- ONLY A: $Z = A * \bar{B}$
- ONLY B: $Z = B * \bar{A}$
- NOT ONLY A: $Z = \overline{(A * \bar{B})}$
- NOT ONLY B: $Z = \overline{(B * \bar{A})}$

Problem 7: Rank order for NRE, Recurring Costs, Flexibility, Performance

Rank order the following design alternatives by filling in the table with 1,2,3,4 representing the relative ranking (1 being the lowest and 4 being the highest). Rank based on the best-case design in each category. If there is a tie, use the lower number (e.g. if tied for 2 or 3, use 2)

	Full-Custom	Std Cell ASIC	Gate Array	FPGA
NRE Costs				
Max Performance				
Energy Efficiency				
Min Per Die Cost				
Flexibility				

Solution:

	Full-Custom	Std Cell ASIC	Gate Array	FPGA
NRE Costs	4	3	2	1
Max Performance	4	3	2	1
Energy Efficiency	4	3	2	1
Min Per Die Cost	1	2	3	4
Flexibility	1	1	2	3

Problem 8: FPGA vs. ASIC

You are the Chief Technical Officer of a new startup, the Citrus Scooter Company. As the chief design lead of your (admittedly small) team of engineers, it is up to you to choose whether to use FPGAs or ASICs for your upcoming line of scooters. After getting quotes from several foundries and having a chat with the Head of HR about the estimated number of work hours needed for each choice, you arrive at the following conclusions for the NRE and per-die costs associated with each approach.

	NRE Cost	Per Die Cost
FPGA	\$25K	\$100
ASIC	\$14.75M	\$5

- As a small startup, you don't anticipate your first deployment to be too large, and the Head of Marketing (who also happens to be Head of HR) estimates that you will deploy about 3000 units in the first quarter. Which type of design would you choose? Why?
- After a successful launch, every city in the country is clamoring for the Citrus scooters! Soon your Head of Marketing comes into your office and excitedly tells you they promised to deploy 200,000 more units by next year. Would you want to switch strategies from your initial approach? If yes, what has now made the other approach so attractive?
- Upon doing some more analyses of the costs involved in switching from one to the other, you find that thankfully your staff have been very organized with their work and some parts can be converted over to the other type of design, which will incur the following NRE costs in converting the design (per-die costs remain the same).

	NRE Cost
FPGA to ASIC	\$12M
ASIC to FPGA	\$8K

How much would it cost to go from your original plan to the other? Is it worth it to switch now that the total cost will be cost of the first run + cost of converting? Is this better than having chosen the other option to begin with?

Solution:

- (a) To find the optimal design strategy, you will have to find the crossover point between the two designs. This is when their total costs are equal.

$$\text{ASIC Cost} = 15,000,000 + N \cdot 5$$

$$\text{FPGA Cost} = 25,000 + N \cdot 100$$

Setting the two equations equal and solving for N :

$$14,750,000 + N \cdot 5 = 25,000 + N \cdot 100$$

$$95N = 14,725,000$$

$$N = 155,000$$

The crossover point is at $N = 155,000$. Since you want to sell less than this amount, FPGA is the more cost efficient strategy.

- (b) With a volume of 200,000 scooters, suddenly ASIC is beginning to look much more attractive, since the per-scooter cost would be significantly lower and as proven in the previous part, this would be enough to offset the NRE from developing an ASIC.
- (c) First let's calculate the cost of the initial FPGA run.

$$\text{Cost}_{\text{FPGA}} = 25000 + (3000)(100) = \$325,000$$

Next, calculate the cost of pivoting from FPGA to ASIC.

$$\text{Cost}_{\text{FPGA2ASIC}} = 12000000 + (200000)(5) = \$13,000,000$$

So together, the total cost will be **\$13,325,000** if we decided on FPGA initially and then changed to ASIC later. What if we had bet everything on ASIC to begin with?

$$\text{Cost}_{\text{ASIC}} = 14750000 + (203000)(5) = \$15,765,000$$

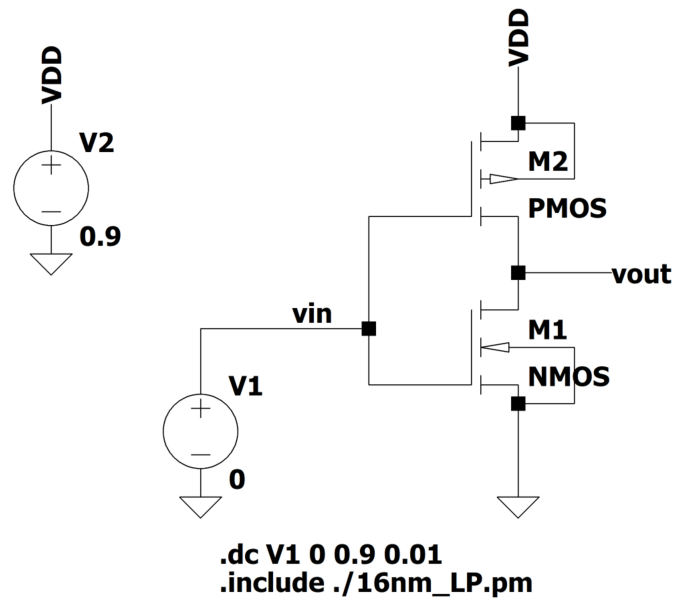
Problem 9: Voltage Transfer Curve (VTC) of Inverter in SPICE

Using SPICE DC Analysis, plot the VTC of an inverter using the 16nm process technology provided on http://ptm.asu.edu/modelcard/LP/16nm_LP.pm. The nominal supply voltage for this process is 0.9V. Use a length of 16 nm for both devices and a width of 3 μm for PMOS and 1 μm for NMOS. Don't forget to connect the body of the devices to the correct voltages - NMOS body to GND and PMOS body to VDD!

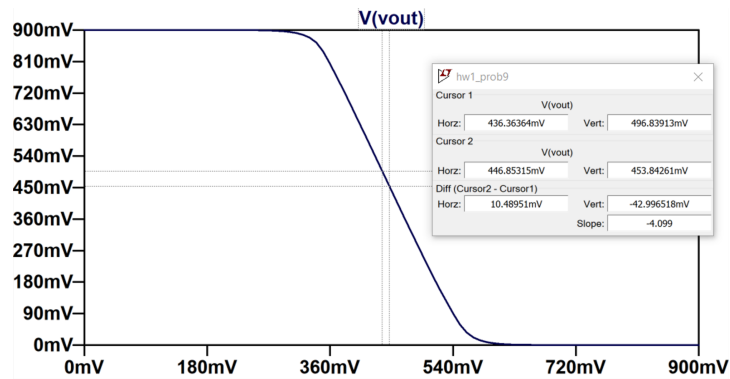
Approximate the maximum voltage gain (slope of $v_{\text{out}}/v_{\text{in}}$). For help setting up the simulation, follow the LTspice tutorial on the class website. https://inst.eecs.berkeley.edu/~eecs151/sp20/files/spice_tutorial.pdf

Solution:

Your inverter schematic should look something like the following:



And the generated VTC curve with the given W and L should look the following:



Using two cursors to grab two separate datapoints in the steepest region of the transfer curve, we get a maximum voltage gain of roughly $a_v = 4.1$.