COS 109 Midterm Exam, Fall 2019 Due 5:00 PM, Friday, October 25

I graded this myself. The median was 63 and the quartiles were 50 and 70; the corresponding values last year were 54, 44 and 59, which suggests that this year's class is better, or the exam was easier, or both. (The extremes were less extreme at both ends this year.) As always, there were difficulties with binary numbers, how information is represented in bits, and the like. Algorithm complexity remains a shaky concept, and the understanding of the Toy machine was imperfect. Arithmetic errors were common. The three values in each column below are for parts 1, 2 and 3, reading from the bottom.

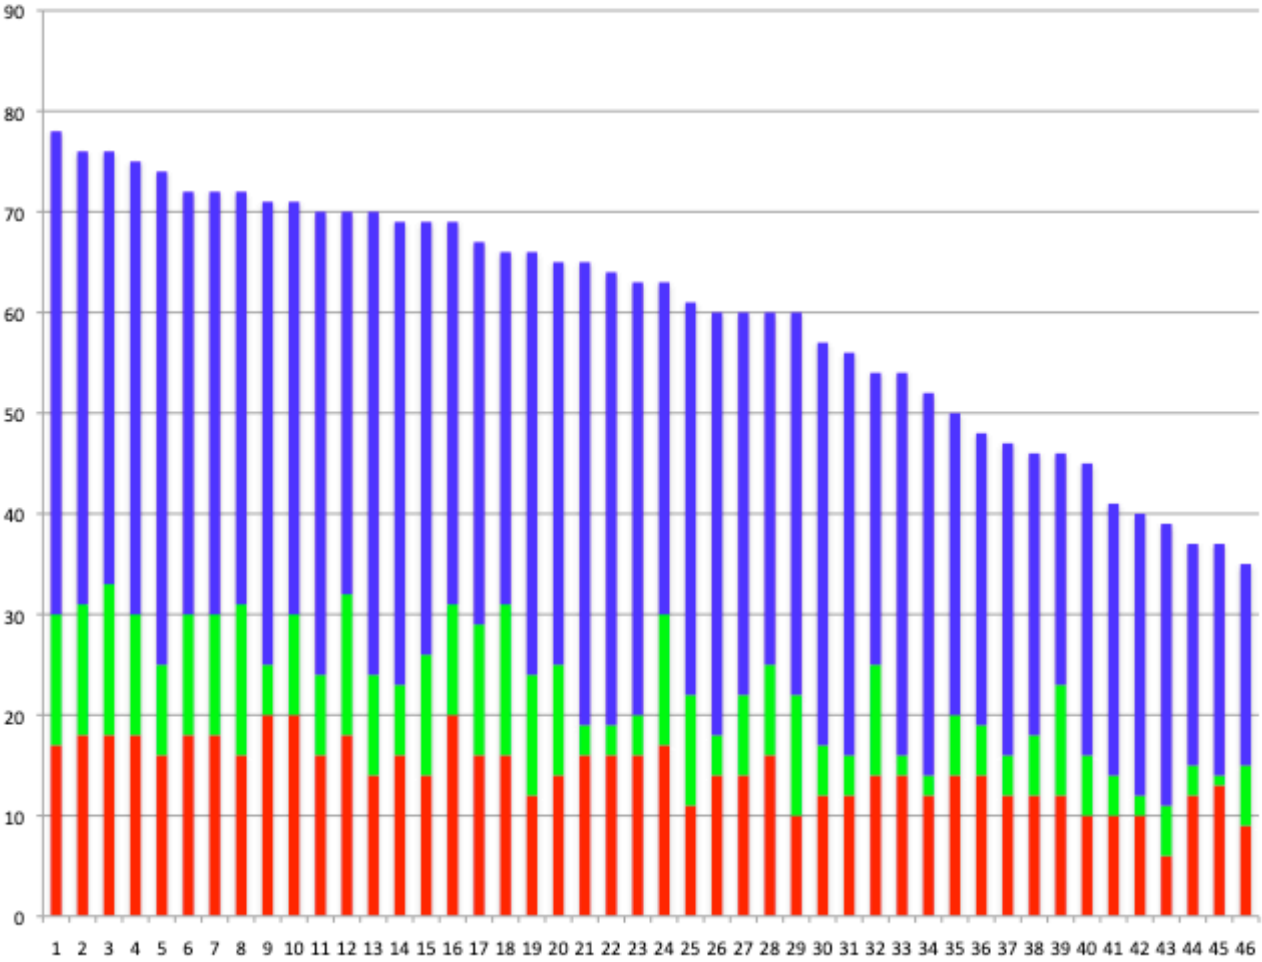

1. (20 points, 2 each) Short Answers. Circle the right answer or write it in the space provided.

(a) Eventually the USA will run out of 9-digit social security numbers. If social security numbers used hexadecimal digits instead of decimal, how many hexadecimal social security numbers could there be? Express your answer as some approximate number multiplied by the closest power of 10, like 1.2 * 10^34.

6.8 * 10^10, or something similar. There are 16^9 numbers.

(b) A while back, the NY Times described a device "the size of a grain of rice" that could be implanted in people as a "subdermal barcode" to help identify them. Which one of the many gadgets and devices we have talked about in class uses the most similar technology?

Prox. Or RFID

(c) In September, a court in ________________________________ decided that the "right to be forgotten" did not apply to Google searches made from outside of the court's jurisdiction. What jurisdiction belongs in the blank?

European Union. Not an individual country, not "Europe," and definitely not in the USA.

(d) The speeds of supercomputers are measured in floating-point operations per second, or "flops." Princeton's web page on 10/7/19 described a new campus supercomputer, Traverse, whose speed of 1.4 PFlops puts it in the top 500 in the world. Approximately how does Traverse's speed compare to the fastest computer on the list? Traverse is…

1/1000 times as fast **1/100 times as fast** **1/10 times as fast** **about the same** **10 times faster**

1/100. The top of the top500 is around 150-200 PFlops.

(e) If **n** and **m** are integers, how many 1 bits (i.e., bits that have the value 1) are there in the binary representation of the number $\mathbf{2^n / 2^m}$?

1. It's a power of 2, regardless of m and n.

(f) If on my laptop I start a big program that I haven't used before, it takes about 10 seconds before it's ready to use. If I then quit the program and 15 seconds later restart it, the startup is much faster. ***Very briefly*** explain why.

Caching. The instructions of the program are likely to still be in the RAM.

(g) As integrated circuit chips get more complicated, they usually need more pins to connect to external circuitry. If pins are arranged around the periphery of a square chip at a fixed spacing as shown in the pictures, which of these expressions best describes how the number of pins varies in proportion to the width **w** of the chip?

log w **√w** **w** **w log w** $\mathbf{w^2}$ $\mathbf{2^w}$

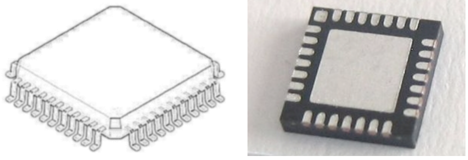

w. The number of pins is proportional to the periphery, which is linear. Not well handled; w^2 and 2^w were popular wrong answers.

(h) Suppose we organize an enormous singles tennis tournament with everyone in the world participating. The winner of each match advances to the next round; the loser of each round is eliminated. (An odd person out gets a bye to the next round.) About how many rounds would it take to determine the ultimate winner?

33. 8 billion people, which is somewhat less than 2^33. Most people got this freebie, which was mentioned in class and probably in the Q/A.

(i) Sports teams often perform a post-game ritual in which each member of one team shakes hands with each member of the other team. If there are **N** players on each team, how many handshakes are there??

N^2. Lots of fancy formulas, not necessary, but full credit if they were correct. Almost this exact question was reviewed in the Q/A.

(j) If I want to store the contents of all of the RAM on the laptops of the 47 students in COS 109 on a single disk, which of these is the smallest that would be sufficient?

1 GB **1 TB** **1 PB** **1 EB** **1 ZB** **1 YB** **none of these**

1 TB. 47 * 8-16GB. Almost everyone got this gift.

2. (15 points) Playing with Toys

Here is a short program in assembly language for the toy computer used in class, with reminders about what the instructions do:

```
Foo  GET          get a number from keyboard into accumulator
     IFZERO Bar   if accumulator is zero, go to location Bar
     IFPOS  Foo   if accumulator is >= zero, go to location Foo
     LOAD   Sum   load accumulator with value in location Sum
     ADD    1     add 1 to accumulator
     STORE  Sum   store accumulator in location Sum
     GOTO   Foo   go to location Foo
Bar  LOAD   Sum
     PRINT
     STOP
Sum  0   reserve a memory location called Sum, set its initial value to 0
```

(a) If you run this program and give it the sequence of inputs **6 5 –3 2 1 –4 7 0** ***exactly*** what does it print?

2. It's counting the negative numbers, easiest seen by just walking through the steps.

(b) Which of these expressions best describes how the running time of this program depends on **N**, the input number?

log N **N** **N log N** **N^2** **2^N** **independent of N**

N. Once around the loop for each input number. The question was not as clearly posed as it should have been; I meant to say "on N, the number of input numbers." I think most people assumed that.

(c) The line "`Sum 0`" could be moved to one other place in this program and the program would behave identically. Where is that place?

After GOTO Foo. Not to the beginning, since the Toy machine assumes that the first thing is a valid instruction.

(d) There are 10 instructions in the Toy computer's repertoire. How many bits are needed to encode an instruction?

4. 2^3 < 10 < 2^4.

(e) Von Neumann's paper describes a computer with 21 instructions. If the Toy computer is modified to include any new instructions from von Neumann's design, could the modified Toy perform fundamentally new computations that it could not have done before? Why do you say so? **Be brief**; we're looking for the idea, not an essay.

No. All computers have the same computational power. I was hoping for mention of "Turing", but didn't see his name very often.

3. (55 points, 5 each) Miscellaneous

(a) Here are five hexadecimal numbers corresponding to the traditional colors Princeton orange, Harvard crimson, Yale blue, Dartmouth green and Podunk gray. Write the proper school name beside each color.

00693E	**Dartmouth (green is dominant)**
0F4D92	**Yale (blue is dominant)**
706E71	**Podunk (all values quite similar)**
C90016	**Harvard (red is dominant)**
E77500	**Princeton (lots of red, significant green)**

Pretty well handled on average.

(b) An article in the *NY Times* some years ago said that the number of regular wired telephone lines in the USA shrank by a factor of two in the previous 12 years, leaving only 65 million wired lines when the story was written.

(i) What is the yearly percentage rate of decline, assuming that it has been a uniform exponential decline?

6%. Factor of two in 12 years: use the rule of 72. Don't use a calculator, and especially don't quote its results to 3 or 4 "significant" figures.

(ii) Assuming (improbably) that this rate of exponential decline continues smoothly, how many years from now will there be only a million wired lines?

72. It takes 6 divisions by two to take 65 million down to 1 million, and thus 6 12-year periods.

(c) Suppose that Thomas Sweet has a special on ice cream cones: they will double the diameter of the scoop for only 4 times the price. Is this a good value for an ice-cream lover, a bad value, or not special at all. Explain your answer by quantitative reasoning.

Good deal. The amount of ice cream is proportional to the cube of the scoop diameter. Doubling the diameter means 8 times the ice cream for 4 times the cost. Many people used area instead of volume, which would be sound reasoning for pizza but not ice cream.

(d) A specific byte in the RAM of a computer contains the hexadecimal value `0F`. Clearly mark each of the following that this single byte could possibly represent.

The letter F in ASCII	**could represent**	[✓ **could not represent**]
Part of an instruction in the Safari browser	[✓ **could represent**]	**could not represent**
An emoji in Unicode	**could represent**	[✓ **could not represent**]
The integer value 15	[✓ **could represent**]	**could not represent**
Part of Allegri's *Miserere* in MP3 format	[✓ **could represent**]	**could not represent**

There is an ASCII table at the end that would confirm that F is something different. Emojis are way up in the Unicode code charts.

(e) According to news stories in October 2019, cameras operated by Los Angeles County police capture and store 3 million license plate pictures every week. The data includes the plate number and the time, date and place where it was recorded. Estimate how much disk space one year's worth of these records will occupy. Explain your reasoning carefully but concisely so we can see how you got your answer.

5-10 GB? Something like 50 bytes to store one item, then multiply by 3 million times 50 weeks in a year. A disturbingly large number of people confused bits and bytes. This question was not as clearly worded as it could have been; I did not mean to include images but full credit if you did and got the numbers and arithmetic right. Reduced credit for excess precision; "3 million" is clearly an approximate figure so you can't have much more than one digit of significance in the result.

(f) Some binary arithmetic:

(i) Add these two binary numbers:

```
11010.101101
  101.010011
------------
```

`100000.` Don't convert to decimal and back! It takes forever and it's easy to get it wrong, as several did.

(ii) Suppose you are adding two **n**-bit binary numbers by hand. Which of these expressions best describes how the amount of work you have to do depends on **n**, the length of each number?

log n **n** **n log n** n^2 2^n **doesn't depend on n**

n. If the numbers were twice as long, it would take twice as long. Not well handled.

(g) Joe College has 100 files on his computer whose names end in `.docx`.

(i) How many times does Joe have to run Word to compute the total number of bytes in all of the `.docx` files?

0. The size information is in the directories.

(ii) How many files does he have to read to determine whether they really are Word files?

100. The filename doesn't have to tell the truth.

This topic was done to death in the Q/A; I hope that those who attended recalled the answers.

(h) Quickies:

Alan Turing's estate endowed the ACM Turing Award	**true**	[✓ **false**]
John von Neumann is buried in his native Hungary	**true**	[✓ **false**]
Bill Gates got his start by writing a Basic interpreter for Windows	**true**	[✓ **false**]
Oracle v Google is about Java APIs	[✓ **true**]	**false**
A prox card is powered by a tiny embedded battery	**true**	[✓ **false**]

I showed pictures of von Neumann's grave in Princeton cemetery. Microsoft was founded 20+ years before Windows came along; you were meant to remember the discussion of the original Altair computer with input via front-panel switches and output via blinking lights.

(i) How does your phone produce a vibration for vibrate mode? Describe briefly but precisely the mechanism by which a vibration is probably produced. 10-15 words should be more than enough.

Motor with off-center/unbalanced weight. Partial credit for saying "motor," smaller credit for "transducer." I passed a opened-up phone around in three separate lectures, each time exhorting everyone to look at it carefully and figure out how it produced vibration. Apparently that wasn't encouragement enough; only a small minority gave the full answer.

(j) The professor in a class with **N** students normally returns problem sets that he has laboriously sorted by student name. For each of the following, give a single expression in **N** (e.g., $2^{\mathbf{N}}$) that tells how the work is proportional to or depends on the size of the class in the worst case.

– If the professor uses an efficient algorithm, how much work does he have to do to sort the problem sets?

N log N. Use Quicksort.

– How many problem sets does the first student have to look at to find her problem set in the sorted pile, if she uses an efficient algorithm?

log N. Use binary search on the sorted pile

– How many problem sets in total must be looked at by all the members of the class when the pile is sorted, if each in turn uses an efficient algorithm to find his or her own problem set?

N log N. N students each do log N work.

– If the professor fails to sort the problem sets, how many problem sets does the first student now have to look at to find her problem set in the unsorted pile?

N. Linear search through the pile.

– How many problem sets in total must be looked at by all members of the class when the pile is unsorted?

N^2. N students each do N work.

No need for complicated formulas, and no need for constant factors like dividing by 2.

(k) ***Exactly*** what do the alphabytes in Jason's cereal bowl say? Write your answer clearly and ***unambiguously***.

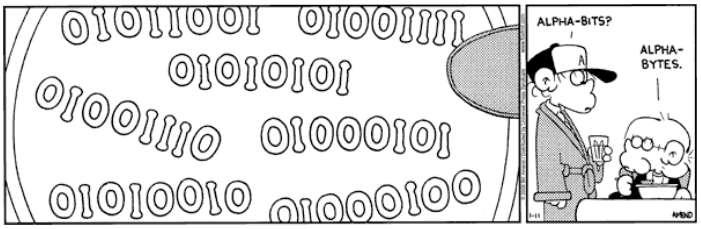

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	NUL	SOH	STX	ETX	EOT	ENQ	ACK	BEL	BS	HT	LF	VT	FF	CR	SO	SI
1	DLE	DC1	DC2	DC3	DC4	NAK	SYN	ETB	CAN	EM	SUB	ESC	FS	GS	RS	US
2	SPC	!	"	#	$	%	&	'	(	)	*	+	,	-	.	/
3	0	1	2	3	4	5	6	7	8	9	:	;	<	=	>	?
4	@	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
5	P	Q	R	S	T	U	V	W	X	Y	Z	[	\	]	^	_
6	`	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o
7	p	q	r	s	t	u	v	w	x	y	z	{	|	}	~	DEL

YOUNERD. The underlined bold italic word "***Exactly***" was meant to imply that an exact answer was required. That means all upper case and no spaces. And "***unambiguously***" was meant to imply clear enough writing that I could be sure of what you had in mind.