

COS 109 Final Exam, Fall 2023

December, 2023

3 hours 180 points total

You may do this exam in any 3-hour period. When you are finished, return the exam to room 311 in the CS building, or scan it (or just the answers) and email it to bwk@cs.princeton.edu. Make sure that any scan is legible! You must return it by 5:00pm EST on Thursday December 21. I would appreciate it if you could tell me how you plan to return it so I know where and when to look, but it's not required.

Please PRINT your name here _____

Honor Pledge: "I pledge my honor that I have not violated the Honor Code during this examination."

Please write the pledge in full and sign it:

This examination is open-book and open-note:

- You may use the textbook, course notes, your own notes, corrected problem sets and solutions, old exams and answer sheets from the course web page, lab instructions, etc.
- You may use a calculator.
- No Internet, Toy simulator or Python compiler is allowed. But you may use a computer, phone or tablet to view course notes, answer sheets, etc., and to write your answers.

There are 180 points for the questions; use the point values for each question to allocate your time (one point per minute). If you're writing or calculating a lot on a question, you may be off on the wrong track.

Write your answers directly on these pages; use the back if necessary. In general, be brief, but if you need more space, attach extra pages and make sure your name is on every extra page. Please write legibly -- I can't grade it if I can't read it. [You can submit just the answers if that makes scanning easier, though showing your work can help with part credit.]

Good luck.

1. (50 pts)

2. (30 pts)

3. (100 pts)

Total

	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

1. (50 points, 2 each) Short Answers. Circle the right answer or write it in the space provided.

(a) I'm thinking of a floating-point number between 1 and 1,000. About how many guesses would it take you to figure it out to within 3 decimal places?

(b) Shrdlu! If I apply a static Huffman coding algorithm to this string of words, it will not compact it as tightly as it might on a not so statistically unusual bunch of words. What's going on? What's odd about all of this? What do you think?

(c) What is a Frankenpine? (Or, if you prefer, a Frankenpalm?)

(d) If it takes one eon to break a message encrypted with AES-128 using brute force, how many eons will it take to break a message encrypted with AES-256?

- 1 2 128 256 128^2 256^2 2^{128} 2^{256}**

(e) Which of these people were knighted for their contributions to computing? Circle all who were.

Charles Babbage Anthony Babington Tim Berners-Lee Tony Hoare Alan Turing

(f) Each of the 50,000 runners in the New York City marathon has an RFID tag to identify them specifically. How many bits does the tag need to use so that each runner has a unique ID?

(g) The country code for Tuvalu is **.tv**. How many web servers does Yahoo need to install in Tuvalu itself so that it can provide some web service for the domain **yahoo.tv**?

(h) A 2001 article about Moore's Law says that Intel "took three decades to produce a chip that ran at one gigahertz, but only 18 months to double that speed." Was Intel's pace of development unusually slow for the three decades, unusually fast for the last year and a half, or just about what would be expected throughout the entire period?

unusually slow unusually fast just about expected

(i) The Great Cannon is a Chinese cyber-warfare tool that "injects malicious JavaScript into pages served from behind the Great Firewall. These scripts, potentially served to millions of users across the Internet, hijack the users' connections to make multiple requests against the targeted site." What kind of attack is this on targeted sites?

DDoS IoT MITM MS-DOS Trojan horse

(j) Which of these historical personages were victims of a cryptographic failure. Circle all that were.

Auguste Kerckhoffs Isoroku Yamamoto Julius Caesar Mary Queen of Scots Queen Elizabeth I

(k) In an IPv4 address, the first part is the network id and the rest is the host id on that net. If there are N bits in the network id part, what is the maximum number of host ids that could be on that network?

(l) The current US Department of Justice vs Google case deals with which one of these legal issues?

antitrust API copyrights patents spoliation trademarks

(m) If **n** is a positive integer, how many 1-bits (that is, bits whose value is 1) are there in the binary representation of $2^{2^2} \dots n$ times, that is, $2^{2^2} \dots$ to a height of **n** ?

- (n) The Luhn algorithm for error-checking a credit card number or a phone IMEI starts at the rightmost digit, and multiplies successive digits alternately by 1 or 2. If the result is > 9 then subtract 9. Add the resulting digits. For a valid number, the sum must be a multiple of 10. How does the running time of this computation grow in proportion to the number of digits in the original number?

log n n n log n n² n³ 2ⁿ no way to predict

- (o) In November 2023, the _____ issued proposed regulations that would require telecom companies to protect their customers better against SIM-swapping attacks. Which one of these entities did this?

FCC FDA FTC ITU NHTSA SEC TIAA

- (p) From a cyber-security bill in the US House of Representatives: “_____ devices have at least one transducer (sensor or actuator) for interacting directly with the physical world, have at least one network interface, and are not conventional Information Technology devices, such as smartphones and laptops.” Which of these belongs in the blank?

5G Bluetooth Cyber-warfare Drone weapons Integrated circuit
Internet of Things Radio frequency identification Two-factor authentication

- (q) Alice says “A crypto algorithm is much more likely to be secure if everyone knows exactly how it works.” Bob says “Nonsense! The only way to make a cryptographic algorithm really secure is to keep how it works a secret.” Eve says “The real security is in the key you use.” Who is right?

none of them only Alice only Bob only Eve both Alice & Eve both Bob & Eve

- (r) If I use my cellphone camera to make a movie of daytime traffic on Nassau Street in Princeton, it has enough memory for about 20 minutes. If instead I make a movie of the night sky while looking for meteor showers, what will I likely discover about the length of movie I can make?

longer movie at night shorter movie at night about the same no way to predict

- (s) In whose collected works would you most likely find the sequence *1, 2, 3, ..., ut, re, mi, fa, sol, la*?

Babbage Bach Goethe Hertz Leibniz Mahler Newton

- (t) Modern processors like those in current PCs and Macs have multiple “cores,” that is, two or more individual CPUs on a single chip. Assuming that all the potential processing power can be perfectly utilized, how does that processing power increase in proportion to n , the number of CPUs on a chip?

logarithmic linear $n \log n$ quadratic cubic exponential

- (u) Whose picture appears on the United Kingdom’s newish 50 pound note, on the opposite side from Queen Elizabeth II?

- (v) Which pair of these acronyms are most closely related? Circle the two closest.

CSS GCC GPT GPU LLM MD5 NAT NDA

- (w) A technical white paper says “5G phone systems should be able to support 1 million connected devices per square kilometer.” If supported devices were spread around uniformly, how many of them could you fit into a square meter?

- (x) What decimal *integer* is the infinitely long binary number **1101100.1111111111...** closest to?

- (y) If I want to create a new top-level Internet domain called **.bwk**, analogous to **.biz**, **.info**, etc., which one of these would have to authorize its creation?

DNS ICANN ITU registrar root server TLA W3C WIPO

2. (30 points) Understanding Programs

- (a) [10 pts] The following Python code is supposed to print a 3-column table that shows each integer from 1 to 100 inclusive, together with its square and its cube; there should also be a line at the beginning of the table that labels the three columns. Sadly, the program doesn't work. Fix the errors by rewriting the code or clearly showing the changes you would make. (This is a question about correct logic, not syntax, but be clear about indentation so I can tell what you mean.)

```
n = 0
while n < 100:
    print("n      n squared      n cubed")
    print(n, n*n, n*3)
```

- (b) [6 pts] Suppose that the Toy machine is augmented with a new instruction **ABS** that replaces the value in the accumulator by its absolute value. That is, if the accumulator value is negative it becomes positive, and if the accumulator is positive it is unchanged. This program uses the **ABS** instruction, with reminders about what the instructions do.

MORE	GET	<i>get a number from user, place it in accumulator</i>	
	IFZERO	END	<i>if accumulator value is 0, go to END</i>
	IFPOS	MORE	<i>if accumulator value is >= 0, go to MORE</i>
	ABS		<i>replace value in accumulator by its absolute value</i>
	ADD	FOO	<i>add value in location FOO to value in accumulator</i>
	STORE	FOO	<i>store value in accumulator in location FOO</i>
	GOTO	MORE	<i>take next instruction from location MORE</i>
END	LOAD	FOO	<i>load value in location FOO into accumulator</i>
	PRINT		<i>print value in accumulator</i>
	STOP		
FOO	0		

If this program is given the sequence of input numbers **2 1 -7 3 -8 -4 -6 5 9 0**, what does it print?

- (c) [3 pts] The fourth line (**ABS**) in this program could be moved to one other place and the program would produce the same answers. Where is that place?

- (d) [4 pts] The Python function **weird** takes two arguments, a list **A** and a value **x**, and returns an integer. In no more than about a dozen words, state clearly what it computes.

```
def weird(A, x):  
    i = 0  
    while i < len(A):  
        if A[i] == x:  
            return i  
        i = i + 1  
    return -1
```

- (e) [2 pts] What is the value returned by **weird([-2,-1,0,1,2], 0)**? The expression in [brackets] is a list.

- (f) [2 pts] What is the value returned by **weird([0,1,2,4,8,16], 3)**?

- (g) [3 pts] Modify the implementation of **weird** in any non-trivial way that preserves its API and correct operation. (fiddling with spacing or changing the names of variables is too trivial.) You only need to indicate clearly what you would change.

3. (100 points, 5 each) Miscellaneous

(a) Molly White, author of the blog `web3isgoinggreat`, uses the Twitter handle `@molly0xFFF`. Suppose that Joe Green and Susan Black decide to copy Molly's idea.

(i) What would Joe Green use in place of `0xFFF` ?

(ii) What would Susan Black use?

(iii) How many potential choices would Earl Gray have that do not collide with Molly, Joe or Susan?

(b) Alice and Bob, bored out of their minds in a COS 109 lecture in Friend 008, are exchanging messages with each other using Gmail from their laptops and texts from their phones. Alice has an Android and uses AT&T; Bob has an iPhone with Verizon. Alice uses Windows while Bob uses macOS. For each of the following statements, assess its likely accuracy.

Their mail messages will use TCP/IP and HTTPS	likely	unlikely
Their text messages will use the router in the ceiling	likely	unlikely
Their text messages will go through different base stations on campus	likely	unlikely
Their mail addresses will be logged by servers at AT&T and Verizon	likely	unlikely
Their text phone numbers will be logged by servers at Google	likely	unlikely

(c) The US International Trade Commission ruled in October 2023 that some versions of the Apple Watch violated pulse-oximetry patents owned by Masimo Corp. For each of these statements, is it likely to be true or false?

Masimo can prevent Apple from importing these watches into the US	true	false
Masimo can prevent Apple from selling these watches in the US	true	false
Masimo can refuse to license its patents to Apple	true	false
Masimo can prevent other watch makers from using its patented technology	true	false
Apple could invent or buy some technology that does not infringe Masimo's patents	true	false

(d) Suppose that the signal received by a cell phone at a distance of one mile from a base station is 100 milliwatts.

(i) How many milliwatts will it receive at a distance of two miles?

(ii) How many milliwatts will it receive at a distance of five miles?

(iii) Why are they called “cell phones”? Mark the right answer.

A biological metaphor, of cells communicating with each other and proliferating as in biological systems

A geometrical metaphor, of honeycomb-like cells that fill a given geographical area

A security metaphor, of locked-down systems that must be jail-broken to access more services

None of these

(e) *Ars Technica* says “OpenAI estimates that it took more than 300 billion trillion floating point calculations to train GPT-3. That’s months of work for dozens of high-end computer chips.” **Very roughly**, how many months would that same computation have taken on your laptop? (A month is about 3 million seconds.)

(f) The Tiobe Index, a somewhat flaky measure of language popularity, reports this long-term history:

2023	2018	2013	2008	2003	1998	1993	1988
1	4	8	6	11	26	19	-
2	2	1	2	2	1	1	1
3	3	4	3	3	2	2	4
4	1	2	1	1	19	-	-

What are the top four languages of 2023, in order? Hint: none are primarily aimed at the web.

(g) At the end of the New York City marathon, the organizers report each runner’s name, place, and finish time, in order of finish time.

(i) If you know a friend’s finish time and you want to find out from this list, as efficiently as possible, what her place was, how long would it take in proportion to n , the number of runners?

- $\log n$
- n
- $n \log n$
- n^2
- 2^n

(ii) If you don’t know anything about your friend’s time, how long will it take?

- $\log n$
- n
- $n \log n$
- n^2
- 2^n

(h) A recent story at Statista.com predicts that the number of IoT devices will grow from 15 billion in 2020 to 30 billion in 2030.

(i) Assuming that this is a smooth exponential growth, what is the *approximate* growth rate per year of the number of Internet-connected devices?

(ii) If growth continues at the same rate, in what year will there be 15 trillion connected devices?

- (i) A story in *Advertising Age* says that T-Mobile is suing Lemonade, an insurance company, over the latter's use of the color magenta, which T-Mobile claims it owns.

(i) What kind of intellectual property is at issue here?

(ii) What specific kinds of intellectual property might the eponymously-named lawyers Phosita and Eula be associated with?

Phosita _____

Eula _____

- (j) A pixel is a picture element and a voxel is a volume element. Suppose you wanted to attach tiny probes all over your body to serve as “touchels”, that is, units of touch. (Whether these might be used for sensing or stimulation we will leave to your imagination.) If each touchel is 0.1 inch by 0.1 inch, estimate **very roughly** how many touchels there would be on your body. You can use metric units if you prefer; if so, assume that touchels are 1 mm by 1 mm. *You must reason quantitatively.* Be sure to state your assumptions clearly.

- (k) Somewhat surprisingly, with 23 people in a room the odds are about 50% that two people will have the same birthday; with 50 people, it's 99.99%. Suppose that each person in the class writes their name and birthday on an index card.

(i) Describe an **efficient** algorithm to determine whether any two people in the class have the same birthday. (Don't worry about multiple duplicates or triples.) **Be clear but brief**; two or three short sentences is enough.

(ii) If there are N people, how does the time that your algorithm takes vary in proportion to N ?

- (l) The US postal service encodes address information in “Intelligent Mail” barcodes like the one in the picture below. There are 65 vertical bars.



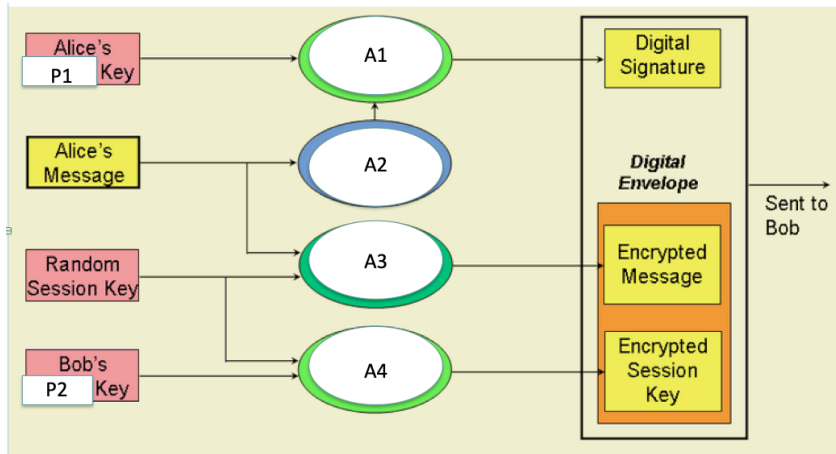
(i) How many different possible addresses could this encoding represent? Just give an expression.

(ii) What is the closest power of 10 to this number?

- (m) In the diagram below, various cryptographic algorithms and terms are used to describe the process of encrypting and digitally signing a message from Alice to Bob. What are valid words or terms to insert in the spaces marked P1, P2, A1, A2, A3, and A4?

P1 _____ P2 _____


A1 _____ A2 _____ A3 _____ A4 _____



- (n) The *NY Times* described a system that allows parents to monitor teenagers who are driving the family car, by visiting a web site that displays the current car location on a map. Assuming that the system is implemented in the most technically feasible and sensible way, assess the likelihood of the following statements:

The position of the car could be monitored by satellite imaging	likely	unlikely
A GPS receiver in the car could broadcast the car’s location to a GPS satellite	likely	unlikely
A GPS-enabled cell phone in the car could report its location to a cell phone base station	likely	unlikely
A cell phone could only report its location when a conversation is in progress	likely	unlikely
RFID would be a viable alternative to a cell phone-based location system	likely	unlikely

- (o) A particular computer network is organized as a *balanced binary tree*: the root computer **R** is directly connected to two child computers, each of which is in turn connected directly to two other child computers, and so on, with no duplicates. “Balanced” means that the connections are made so that as much as possible each computer has exactly two children. For instance, if another computer were added to the network below, it would be added as the right child of **c2**.



- (i) Which of these terms best describes how the number of connecting wires will grow in proportion to **n**, the number of computers on the network?

logarithmic linear $n \log n$ quadratic cubic exponential

- (ii) Which of these terms best describes how the maximum distance from any computer to any other computer in the network will grow in proportion to **n**?

logarithmic linear $n \log n$ quadratic cubic exponential


- (p) UTF-8 is a variable-length encoding of Unicode that represents ASCII characters as one byte and other characters as two or more bytes. It uses one or more of the leftmost bits of each byte to indicate the length of the character: 1-byte characters are represented as `0xxxxxxx`, and 2-byte characters are represented as `110xxxxx 10xxxxxx`, where each `x` is either 0 or 1. (This encoding saves space if the text is mostly ASCII.)

- (i) How many one-byte characters can this format encode? (An expression is ok.)

- (ii) How many two-byte characters can this format encode? (An expression is ok.)

(q) The picture on the left shows a pattern of bricks protruding from the wall of a campus building. I've drawn it more clearly on the right.

(i) *Exactly* what does the pattern say?



(ii) Bonus: What is the answer? **Yes** **No** **We don't know**

(r) An IPv4 address is a 32-bit integer.

(i) If the IPv4 address **63.254.255.255** is stored in a 32-bit integer variable **v** and incremented by the statement **v = v+1**, what is the resulting value, also expressed in dotted decimal notation?

(ii) What is that resulting value expressed in hexadecimal?

(s) [10 pts] Quickies. Circle the best answers:

- | | | |
|--|-------------|--------------|
| Fei-Fei Li '99 won a Turing award for her development of ImageNet | true | false |
| “Satellites track your cellphones and can tell 911 operators where you are” | true | false |
| Intel dominates Nvidia in the GPU marketplace | true | false |
| A two-factor device is hardware used to do efficient prime testing for the RSA algorithm | true | false |
| In supervised learning, a human evaluates each decision made by an ML system | true | false |
| An ML word vector encodes directional information like “north” and “south” | true | false |
| Bitcoin transactions are anonymous to the US Internal Revenue Service | true | false |
| “AI Winter” refers to an artificial intelligence funding cycle that ends every December | true | false |
| It is possible for there to be more static web pages than IPv4 addresses | true | false |
| Javascript code can easily monitor where you move your mouse on your laptop screen | true | false |