

COS 226	Algorithms and Data Structures	Fall 2014
Final		

This test has 14 questions worth a total of 100 points. You have 180 minutes. The exam is closed book, with the exception of a one page cheatsheet. No calculators or other electronic devices are permitted. **Write out and sign the Honor Code pledge just before turning in the test.**

This exam is preprocessed by computer. Please use a pen; if you use a pencil, be sure to write darkly. Do not write any answers outside of the designated frames. And do not write on the corners.

“I pledge my honor that I have not violated the Honor Code during this examination.”

Name:

netID:

Room:

Precept: P01 P02 P03 P03A P04 P04A

Problem	Score
0	
1	
2	
3	
4	
5	
6	
Sub 1	

Problem	Score
7	
8	
9	
10	
11	
12	
13	
Sub 2	

P01	F 9	Andy Guna
P02	F 10	Jérémie Lumbroso
P03	F 11	Josh Wetzel
P03A	F 11	Jérémie Lumbroso
P04	F 12:30	Robert MacDavid
P04A	F 13:30	Shivam Agarwal


Total	
-------	--

0. Initialization (2 points)

In the space provided on the front of the exam, write your name and Princeton netID; fill in your precept number; write the name of the room in which you are taking the exam; and write and sign the honor code.

1. Digraph Traversal (6 points)

Consider the following digraph. Assume the adjacency lists are in sorted order: for example, when iterating through the edges pointing from vertex 5, consider the edge $5 \rightarrow 3$ before the others.



(a) Starting from vertex 0, run a depth-first search of the digraph, and list the vertices in *reverse postorder*.

--	--	--	--	--	--	--	--	--

(b) Starting from vertex 0, run a depth-first search of the digraph, and list the vertices in *preorder*.

--	--	--	--	--	--	--	--	--

3. String Sorting Algorithms (7 points)

The column on the left is the original input of 24 strings to be sorted; the column on the right are the strings in sorted order; the other 7 columns are the contents at some intermediate step during one of the 3 radix sorting algorithms listed below.

Match up each column with the corresponding sorting algorithm. You may use a number more than once.

Hint: think about algorithm invariants; do not trace code.

leaf	cost	hash	edge	rank	load	find	cost	cost
size	edge	edge	cost	hash	leaf	load	edge	edge
null	flow	cost	fifo	edge	heap	size	find	fifo
type	find	heap	flow	leaf	swap	type	fifo	find
cost	fifo	fifo	find	heap	node	trie	flow	flow
sink	heap	flow	heap	less	fifo	node	heap	hash
heap	hash	find	hash	next	edge	edge	hash	heap
trie	leaf	leaf	leaf	fifo	trie	time	leaf	leaf
loop	loop	loop	loop	time	swim	leaf	loop	less
flow	less	load	load	find	null	push	less	list
less	load	less	less	sink	time	hash	load	load
node	list	list	list	list	find	sink	list	loop
find	null	next	next	size	sink	rank	null	next
next	node	node	node	flow	rank	null	node	node
fifo	next	push	push	load	loop	swim	next	null
push	push	rank	rank	node	flow	fifo	push	push
rank	rank	trie	trie	loop	type	heap	rank	rank
load	size	sink	sink	cost	push	loop	size	sink
edge	sink	type	type	trie	hash	swap	sink	size
hash	swap	time	time	null	less	less	swap	swap
time	swim	swap	swap	push	cost	cost	swim	swim
swap	type	null	null	swap	list	next	type	time
list	trie	swim	swim	swim	next	list	trie	trie
swim	time	size	size	type	size	flow	time	type

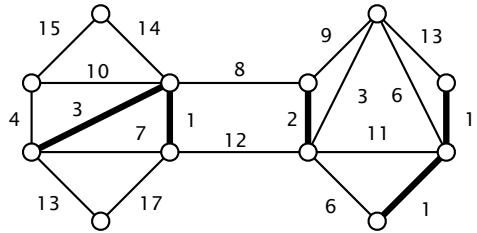
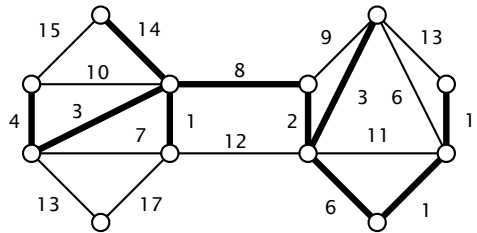
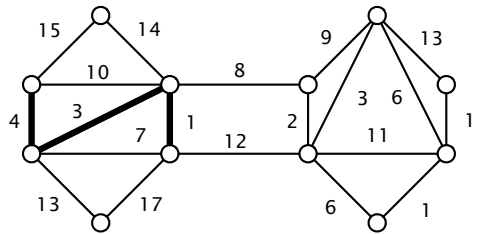
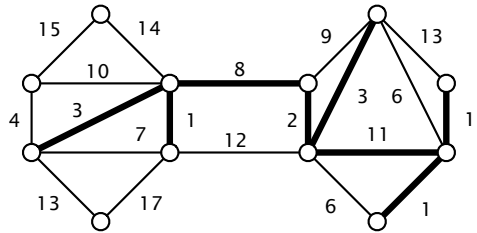
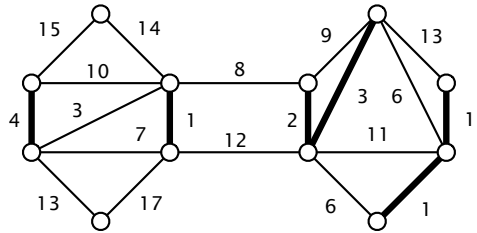
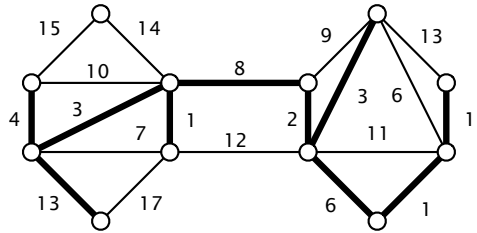
0

4

- (0) Original input
- (1) LSD radix sort
- (2) MSD radix sort
- (3) 3-way radix quicksort (no shuffle)
- (4) Sorted

5. Minimum Spanning Tree Algorithms (6 points)

Each of the figures below represents a partial spanning tree. Determine whether it could possibly be obtained from (a prematurely stopped) Prim's algorithm, (a prematurely stopped) Kruskal's algorithm, both or neither.

	PRIM	KRUSKAL	BOTH	NEITHER
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

7. Properties of Algorithms (9 points)

Check whether each of the following statements are True or False.

- (a) **Shortest paths.** Consider an edge-weighted digraph G with distinct and positive edge weights, a source vertex s , and a destination vertex t . Assume that G contains at least 3 vertices, has no parallel edges or self loops, and that every vertex is reachable from s .

	True	False
Any shortest $s \rightarrow t$ path must include the lightest edge.	<input type="radio"/>	<input type="radio"/>
Any shortest $s \rightarrow t$ path must include the second lightest edge.	<input type="radio"/>	<input type="radio"/>
Any shortest $s \rightarrow t$ path must exclude the heaviest edge.	<input type="radio"/>	<input type="radio"/>
The shortest $s \rightarrow t$ path is unique.	<input type="radio"/>	<input type="radio"/>

- (b) **Minimum spanning trees.** Consider an edge-weighted graph G with distinct and positive edge weights. Assume that G contains at least 3 vertices, has no parallel edges or self loops, and is connected.

	True	False
Any MST must include the lightest edge.	<input type="radio"/>	<input type="radio"/>
Any MST must include the second lightest edge.	<input type="radio"/>	<input type="radio"/>
Any MST must exclude the heaviest edge.	<input type="radio"/>	<input type="radio"/>
The MST is unique.	<input type="radio"/>	<input type="radio"/>





- (c) **Burrows-Wheeler transform.**

	True	False
Any input x consisting of an integer (between 0 and $N - 1$) followed by N characters is the Burrows-Wheeler transform of some string s of length N .	<input type="radio"/>	<input type="radio"/>
If the Burrows-Wheeler transforms of s and t are equal, then $s = t$.	<input type="radio"/>	<input type="radio"/>
If the Burrows-Wheeler inverse transforms of x and y are equal, then $x = y$.	<input type="radio"/>	<input type="radio"/>
In practice, applying the Burrows-Wheeler transform is significantly faster than applying the Burrows-Wheeler inverse transform.	<input type="radio"/>	<input type="radio"/>

8. **Huffman Trees** (4 points)

Consider the string “DATA-STRUCTURES-AND-ALGORITHMS”: which of the following trees is an optimal prefix-free code for this input string?

Optimal Prefix-Free Code Not Optimal Prefix-Free Code



9. **LZW Compression** (5 points)

What is the result of compressing the following string of length 15 using LZW compression?

B B B B B B C A B B C B B B C

Assume the original encoding table consists of all 7-bit ASCII characters and uses 8-bit codewords. Recall that codeword 80 is reserved to signify end of file.

42									80
----	--	--	--	--	--	--	--	--	----

For reference, below is the hexadecimal-to-ASCII conversion table from the textbook:

	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	SP	!	"	#	\$	%	&	'	()	*	+	,	-	.	/
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

10. Burrows-Wheeler Transform (6 points)

(a) What is the Burrows-Wheeler transform of the following?

A D D B D B C A

(b) What is the Burrows-Wheeler inverse transform of the following?

2
C A D D A B C C

--	--	--	--	--	--	--	--	--	--

Feel free to use both of these grids for scratch work.

11. Algorithm and Data Structure Design (13 points)

Design a data type to store a collection of gene fragments over the DNA alphabet $\{A, C, T, G\}$, according to the following API:

```
public class FragmentCollection


---


public      FragmentCollection()    create an empty collection of DNA fragments
public void  add(String fragment)   add the DNA fragment to the collection
public int   prefixCount(String p)  number of DNA fragments that start with prefix p
```

Here is an example:

```
FragmentCollection fc = new FragmentCollection();
fc.add("AC");
fc.add("TACG");
fc.add("TCGAA");
fc.add("CGA");
fc.add("AGCT");
fc.add("TCGG");
fc.add("TCGG");           // added twice, will be counted twice
fc.prefixCount("");      // returns 7 (number of adds)
fc.prefixCount("T");     // returns 4 (TACG, TCGAA, TCGG, TCGG)
fc.prefixCount("TC");    // returns 3 (TCGAA, TCGG, TCGG)
fc.prefixCount("G");     // returns 0
```

Give a crisp and concise English description of your data structure. Your answer will be graded on correctness, efficiency, and clarity.

- (a) Declare the instance variables for your `FragmentCollection` data type. You may use nested data types.

```
public class FragmentCollection {
```


```
}
```


12. Reductions (13 points)

Consider the following two graph-processing problems:

- **SHORTEST-PATH.** Given an edge-weighted digraph G with nonnegative edge weights, a source vertex s and a destination vertex t , find a shortest path from s to t .
- **SHORTEST-TELEPORT-PATH.** Given an edge-weighted digraph G with nonnegative edge weights, a source vertex s and a destination vertex t , find a shortest path from s to t where you are permitted to *teleport* across one edge for free. That is, the weight of a path is the sum of the weights of all of the edges in the path, excluding the largest one.

For example, in the edge-weighted digraph below, the shortest path from s to t is $s \rightarrow w \rightarrow t$ (with weight 11) but the the shortest teleport path is $s \rightarrow u \rightarrow v \rightarrow t$ (with weight 3).



- (a) Design a linear-time reduction from **SHORTEST-PATH** to **SHORTEST-TELEPORT-PATH**. To demonstrate your reduction, draw the edge-weighted digraph (and label the source and destination vertices) that you would construct to solve the **SHORTEST-PATH** problem on the digraph above. You may additionally explain your construction with a few concise sentences.

- (b) Design a linear-time reduction from SHORTEST-TELEPORT-PATH to SHORTEST-PATH. To demonstrate your reduction, draw the edge-weighted digraph (and label the source and destination vertices) that you would construct to solve the SHORTEST-TELEPORT-PATH problem on the digraph given in the previous page. You may additionally explain your construction with a few concise sentences.

- (c) Determine whether each of following statements can be inferred from the fact that SHORTEST-PATH and SHORTEST-TELEPORT-PATH linear-time reduces to one another. For simplicity, assume $E \geq V$.

	Yes	No
If there exists an $E \log \log E$ algorithm for SHORTEST-TELEPORT-PATH, then there exists an $E \log \log E$ algorithm for SHORTEST-PATH.	<input type="radio"/>	<input type="radio"/>
If there exists an $E \log \log E$ algorithm for SHORTEST-PATH, then there exists an $E \log \log E$ algorithm for SHORTEST-TELEPORT-PATH.	<input type="radio"/>	<input type="radio"/>
If there does not exist a linear-time algorithm for SHORTEST-PATH, then there does not exist a linear-time algorithm for SHORTEST-TELEPORT-PATH.	<input type="radio"/>	<input type="radio"/>
If there does not exist a linear-time algorithm for SHORTEST-TELEPORT-PATH, then there does not exist a linear-time algorithm for SHORTEST-PATH.	<input type="radio"/>	<input type="radio"/>

13. Problem Identification (9 points)

You are applying for a job at a new software technology company. Your interviewer asks you to identify the following tasks as either *possible* (with algorithms and data structures introduced in this course), *impossible*, or an *open research problem*.

	Possible	Impossible	Open
Given a digraph, find a directed cycle that is simple (if one exists) in time proportional to $E + V$. A <i>simple</i> cycle is a cycle that has no repeated vertices other than the requisite repetition of the first and last vertex.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Given an edge-weighted digraph in which all edge weights are either 1 or 2 and two vertices s and t , find a shortest path from s to t in time proportional to $E + V$.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Given an edge-weighted DAG with positive edge weights and two vertices s and t , find a <i>path</i> from s to t that <i>maximizes the product</i> of the weights of the edges participating in the path in time proportional to $E + V$.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Given an edge-weighted graph with positive edge weights, find a <i>spanning tree</i> that <i>maximizes the product</i> of the weights of the edges participating in the spanning tree in time proportional to $E + V$.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Given an edge-weighted graph with positive edge weights and two distinguished vertices s and t , find a <i>simple path</i> (no repeated vertices) between s and t that <i>maximizes the sum</i> of the weights of the edges participating in the path in time proportional to EV .	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Given a flow network and a mincut in that flow network, find a maxflow in time proportional to $E + V$.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Given an array of N strings over the DNA alphabet $\{A, C, T, G\}$, determine whether all N strings are distinct in time linear in the number of characters in the input.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Given an array a of N 64-bit integers, determine whether there are two indices i and j such that $a_i + a_j = 0$ in time proportional to N .	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Given an array of N integers between 0 and $R^2 - 1$, <i>stably sort</i> them in time proportional to $N + R$.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>