Qualifying Exam (August 2020): Operations Research

You have 4 hours to do this exam. **Reminder:** This exam is closed notes and closed books.

Do 2 out of problems 1,2,3.
Do 2 out of problems 4,5,6.
Do 3 out of problems 7,8,9,10,11,12,13,14,15,16.

All problems are weighted equally. **On this cover page write which seven problems you want graded.**

problems to be graded:

Academic integrity is expected of all students at all times, whether in the presence or absence of members of the faculty. Understanding this, I declare that I shall not give, use, or receive unauthorized aid in this examination.

Name (PRINT CLEARLY), ID number

Signature

Detailed Instructions for Taking This Exam Over Zoom

1. This exam is conducted **via Zoom on August 19, 2020, from 9:00 am to 1:00 pm EDT**.
2. The entire Zoom meeting and chat messages are being recorded.
3. This is a closed book, closed note exam.
4. Hand calculators (or other computing devices) may not be used during the exam.
5. Zoom: You should join the Zoom meeting from two devices: Your computer/laptop/tablet (with webcam), and your smartphone (with camera).
6. Audio should be muted, and video must be kept on during the exam.
7. Your computer webcam must fully show your face; your smartphone camera should show your hands and workspace, with the pages of paper being used for the exam.
8. At the very beginning of the exam, during set up, you will be asked to do a brief "environment scan", showing the workspace where your computer is and the desk/table/floor where you will be writing your work.
9. You are required to bring enough blank pieces of paper to use for the exam. You will show the blank pages at the beginning, during the "environment scan" on Zoom.
10. If you do not have a printer to be able to print the exam cover sheet, then simply write the critical information from the cover sheet (your name, problems to be graded, and your id number and signature) on a blank piece of paper and use that as your cover sheet. Also, if a figure appears in a question, you can replicate the figure, as needed, as best possible on a blank sheet of paper (possibly holding the page over the computer screen to trace the figure).
11. You are not allowed to use the internet for any searching or communication with others, with the sole exception of communication with the proctor(s) via Zoom chat (which is set so that your chats only go privately to hosts, not to others).
12. Use a clean piece of paper to answer each question. Multiple pages for one problem may be used, as needed. (You should show your work!) Clearly label all pages.

    If the problem has an associated figure (on the exam page), you can write directly on a printout of that figure and include it in the scanned solutions you send by email.
13. Scan and submit your answers in a single pdf file. Make sure you scan carefully, so that images are clear, not blurred or truncated.
14. No students are allowed to leave the Zoom meeting until the exam is over.
15. If you finish the exam early, then submit your exam and remain in the Zoom meeting until the conclusion of the exam at 1:00pm, EDT.
16. After submitting your exam, you can study for another exam or work on anything else, while staying in view in the Zoom meeting.
17. If you have a question during the exam, then send a Zoom chat message (it will go only to the host(s), not to others).
18. After you finish the exam, scan all of your pages of work, into a single pdf, and email as an attachment to `esther.arkin@stonybrook.edu` no later than 5 minutes after completion of the exam (i.e., **by 1:05pm, EDT, on Wed, Aug 19**).

(1). Consider the following two LPs:

$$\begin{cases} (LP1) & \max & \mathbf{c}^T\mathbf{x} \\ & \text{st} & A\mathbf{x} \le \mathbf{b} \\ & & \mathbf{x} \lesseqqgtr 0 \end{cases}$$

$$\begin{cases} (LP2) & \min & \mathbf{c}^T\mathbf{x} \\ & \text{st} & A\mathbf{x} \ge \mathbf{b} \\ & & \mathbf{x} \lesseqqgtr 0 \end{cases}$$

1. Give the duals of the two problems.
2. Suppose both (LP1) and (LP2) are feasible. Prove that if one of them has a finite optimal solution then so does the other.
3. Suppose both (LP1) and (LP2) are feasible. Prove that if one of them has an unbounded objective function then so does the other.
4. Suppose that both (LP1) and (LP2) have finite optimal solutions. Let x_1 be a feasible point to (LP1) and x_2 be a feasible point to (LP2). Prove that $\mathbf{c}^T\mathbf{x}_1 \le \mathbf{c}^T\mathbf{x}_2$.

(2). Two reservoirs are available to supply the water needs of three cities. Each reservoir can supply up to 50 million gallons of water per day (MGD). Each city needs 40 MGD. For each MGD of unmet demand there is a penalty. The costs (in thousands of dollars) of transporting 1 MGD and penalties are given in the following table:

	City 1	City 2	City 3
Reservoir 1	7	8	10
Reservoir 2	9	7	8
Penalty	20	22	23

Formulate a balanced transportation tableau for this problem and find the optimal solution.

(3). Consider the following integer programming problem:

$$\begin{array}{llrcl} \max & z = & x_1 + 4x_2 & & \\ \text{s.t.} & & x_1 + 2x_2 & \le & 7 \\ & & -x_1 + 2x_2 & \le & 3 \\ & & x_1\ ,\ x_2 & \ge & 0\ ;\ x_1\ ,\ x_2 \text{ integer} \end{array}$$

Solve this problem by first finding the LP optimum and then finding the integer optimum by the cutting plane method.

(4). Let $\{N(t), t \ge 0\}$ be a Poisson process with arrival rate λ, and denote by S_n the time of the occurrence of the nth event in $\{N(t)\}$. Given that $N(t) \le 1$ for some fixed time t, find the conditional density function of S_2.

(5). Consider an $M/G/1$ queuing system with Poisson arrivals $PP(\lambda)$ and i.i.d. service times $S_1, S_2, \ldots$. Let $\mu = E[S_1]$ and $\sigma^2 = E[S_1^2]$. Assume that $\lambda\mu < 1$. Use the mean-value approach (i.e., by combining the PASTA property with Little's law) to compute L, the long-run expected number of customers in system (*your answer should only contain* λ, μ, *and* σ^2).

(6). A coin is tossed repeatedly and independently. Suppose the probability the coin turns heads is $p \in (0,1)$. Let $X_n \in \{H, T\}$ be the outcome of the nth toss. Define $Y_n = 1$ if $X_n = H,\ X_{n-1} = T,\ X_{n-2} = H$ for $n \ge 3$; $Y_n = 0$ otherwise. In other words, $Y_n = 1$ is the event that the sequence HTH is observed on the nth toss. Use the discrete-time Markov chain theory to compute the long run probability $\lim_{n\to\infty} P(Y_n = 1)$.

(7). (a). How efficiently (in big-Oh notation) can the (Euclidean) Delaunay diagram be constructed for n points in the plane? What about for n points in 3D? *Describe briefly (without giving details) one method that achieves the time bound in 2D.*

(b). State an important property of the Euclidean Delaunay diagram that might be used in an application. Sketch (briefly) how that property would be proved.

(c). For the set S of 11 points shown below, draw the (Euclidean) Delaunay diagram. In order to assist you in making some decisions, I have drawn some possibly relevant circles. NOTE: I provide 2 copies of the figure, in case you make a mistake on one.

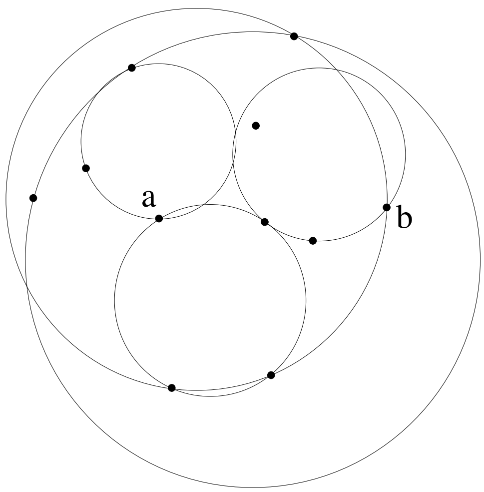

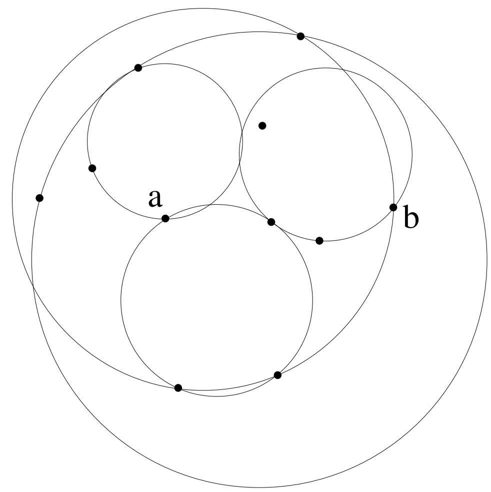

(d). Consider now the furthest-site Delaunay diagram of S.

(i). How efficiently can it be computed (in big-Oh)?

(ii). Explain briefly the relationship between the furthest-site Delaunay diagram and the three-dimensional convex hull of points related to S.

(iii). Below, draw the furthest-site Voronoi cell associated with point b.

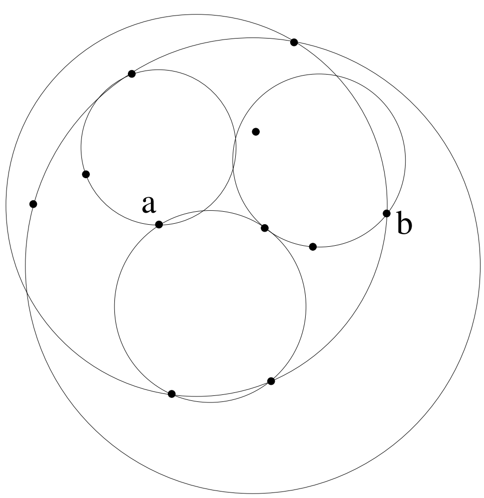

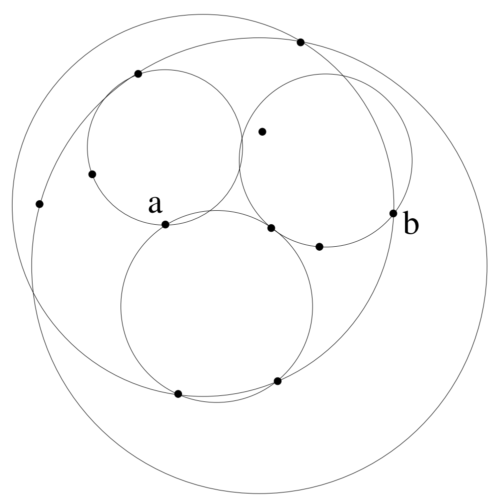

(8). Given a set $S = \{p_1, \ldots, p_n\}$ of n points in the plane. We are to build a data structure to support efficient queries: For a (nonvertical) query line L, find the first point of S to be hit when L is translated upwards (in $+y$ direction). (or report that no point is hit, if all of S lies below L) For example, in the figure below, the result of a query with line L is to report the point p_i, shown in red.

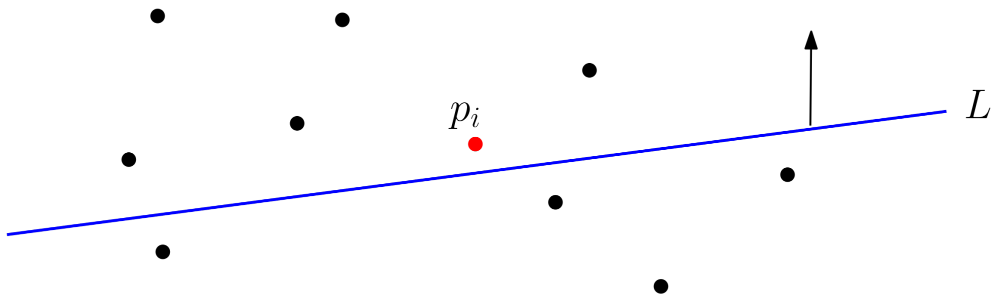

Our goal is to have particularly efficient query time, after doing some preprocessing to construct a data structure (of reasonable size).

(a). What preprocessing/storage/query can be achieved? (Best that you know how to achieve.)

(b). Describe what steps are needed to achieve your answer to part (a). Justify your answer and describe briefly any data structures used.

(9). Let $X_1, \ldots, X_n$ be independent random variables, each with mean 0, and each with finite third moments. Show by using characteristic functions that

$$E\left\{\left(\sum_{i=1}^{n} X_i\right)^3\right\} = \sum_{i=1}^{n} E\{X_i^3\}.$$

(10). For X, Y in L^2 show

$$(E\{XY|\mathcal{G}\})^2 \le E\{X^2|\mathcal{G}\}E\{Y^2|\mathcal{G}\}$$

(the Cauchy-Schwartz inequality).

(11). Let $f(x)$ be a given probability density function. Give the detailed steps of the Acceptance-Rejection method for generating random variates from $f(x)$. Also demonstrate the validity of the method, i.e., show that $P(X \le x) = \int_{-\infty}^{x} f(y)dy$, where X is the random variate produced by the algorithm.

(12). Let X_1 and X_2 be two independent random variables uniformly distributed on $[0,1]$. Define $X = X_1 + X_2$. Provide two algorithms for generating a random variate with the same distribution as X. Compare the efficiency of the two algorithms and determine which one is preferable.

(13). Let us consider a discounted MDP with a state space $X = \{s_1, s_2\}$; action sets $A(s_1) = \{a_{1,1}, a_{1,2}\}$ and $A(s_2) = \{a_{2,1}\}$; one-step rewards $r(s_1, a_{1,1}) = 1$, $r(s_1, a_{1,2}) = 0$, and $r(s_2, a_{2,1}) = 2$; and transition probabilities $p(s_1|s_1, a_{1,1}) = 1$, $p(s_2|s_1, a_{1,2}) = 1$, and $p(s_2|s_2, a_{2,1}) = 1$. Let $v_\alpha^\pi(x) := E_x^\pi \sum_{t=0}^{\infty} \alpha^t r(x_t, a_t)$ be the expected total discounted reward for the discount factor α, policy π, and the initial state $x \in X$. We recall that a nonrandomized policy ϕ is called stationary if the decision at each state depends only on the state. In other words, ϕ is a stationary policy if it chooses the decision $\phi(x)$ at each state $x \in X$.

Show that for $\alpha \in [0, 0.5]$ the stationary policy ϕ with $\phi(s_1) = a_{1,1}$ is optimal, and, if $\alpha \in [0.5, 1)$, the stationary policy ϕ with $\phi(s_1) = a_{1,2}$ is optimal.

(14). For the the MDP described in the previous problem and for two discount factors α_1 and α_2, let us consider the objective criterion $g_{1,2}^\pi(x) = v_{\alpha_1}^\pi(x) + v_{\alpha_2}^\pi(x)$ for a policy π and initial state $x \in X$. Suppose $\alpha_1 = 0.2$ and $\alpha_2 = 0.6$. Let σ be the Markov policy choosing at state s_1 action $a_{1,1}$ for the initial period and then $a_{1,2}$ for all subsequent periods. Show that $g^{\sigma}_{1,2}(x) \geq g^{\phi}_{1,2}(x)$ for all stationary policies ϕ and for all initial states $x \in X$.

(15). Which of the following claims are true and which are false? Justify your answer by giving a short proof or a counterexample:
(a) If all arcs in a network have different costs, the network has a unique shortest path tree.
(b). In a directed network with positive arc lengths, if we eliminate the direction on every arc (i.e., make it undirected) the shortest path distances will not change.
(c). In a shortest path problem, if each arcs length is increased by k units, the shortest path distances increase by a multiple of k.
(d). In a shortest path problem, if each arcs length is decreased by k units, the shortest path distances decrease by a multiple of k.
(e). Among all shortest paths in a network, Dijkstra's algorithm always finds a shortest path with the least number of arcs.

(16). Consider a max flow problem on $G = (N, A)$ with source node s and sink node t, upper bounds on the flows u_{ij} are integral or infinite.
(a). Prove that a maximum flow v is finite if and only if the network does not contain a directed path from s to t whose arc capacities are all infinity.
(b). Suppose that the network has some infinite capacity arcs, but no infinite capacity path from s to t. Let A^0 be the arcs with finite capacity in A. Prove that we can replace the capacities on arcs of $A \setminus A^0$ by a finite number $M \geq \sum_{(i,j)\in A^0} u_{ij}$ without affecting the max flow value v.