

Instructions: Please answer the following questions to the best of your ability. If you are asked to design an algorithm, please describe it clearly and concisely, prove its correctness, and analyze its running time. If you are asked to show your work, please include relevant calculations for deriving your answer. If you are asked to explain your answer, give a short (~ 1 sentence) intuitive description of your answer. If you are asked to prove a result, please write a complete proof at the level of detail and rigor expected in prior CS Theory classes (i.e. 103). When writing proofs, please strive for clarity and brevity (in that order). Cite any sources you reference.

1 (13 points) Poe learns of Karger's Algorithm

- (a) (3 points) Poe hears about Karger's algorithm to solve the minimum cut problem for unweighted graphs. Poe now wants to figure out how to solve the minimum cut problem for *weighted* graphs. Now, the size of a cut is measured as the sum of the weights of the edges going across the cut, and Poe wants to minimize this weight sum over all cuts. Barr has devised a secret algorithm that he refuses to share with Poe. Poe goes to Barr's evil twin brother Burr the Bear. Burr tells Poe that he can solve the problem for weighted graphs by replacing each edge of weight k by a path of k edges of unit weight. That is, if there is an edge (u, v) of weight k , then Poe should delete the edge, rename u by $x_{u,v,0}$ and v by $x_{u,v,k}$, add new vertices $x_{u,v,1}, x_{u,v,2}, \dots, x_{u,v,k-1}$ and add new edges of unit weight between $x_{u,v,i}$ and $x_{u,v,i+1}$ where $i \in [0, k-1]$. After all edges are replaced by paths of unit-weight edges, the resulting graph is an unweighted graph.

Poe can then use Karger's algorithm to solve his problem (assume that Poe runs Karger's algorithm so many times that with probability 1 Poe will find a min-cut in the unweighted graph). The solution will be a minimum cut C (which is a partition of vertices into two non-empty subsets) for the unweighted graph. Poe then remove all vertices $x_{u,v,i}$ for all u, v and $i \in [1, w(u, v) - 1]$ from C to obtain a partition of vertices in the original graph (let us call it C'). Poe hopes that C' is a min-cut in the original, weighted graph, but sadly this scheme fails in many cases.

Find a small counterexample (with at most 5 vertices). (In your example all edge weights should be positive integers.) Your example should show that C' is not a valid cut or is not a min-cut.

- (b) (3 points) The max-cut of a graph is a cut with maximum number of edges crossing the cut. The complement G' of a graph G is defined as a graph on the same set of vertices as G but so that there is an edge (u, v) in G' if and only if there is no edge (u, v) in G (hence complementing the edges). Consider any cut S of G and the same cut S in G' . Poe uses the following logic: the fewer edges there are across S in G , the more non-edges there across S in G and hence the more edges there are across S in G' . Because of this, Poe thinks that given a max-cut of some graph G , it should correspond to a min-cut of the complement of G .

Formally, let $c_{\min}(G)$ be the size of a min-cut in graph G and let $c_{\max}(G)$ be the size of a max-cut in G . If C is a min-cut in graph G (that is, C has $c_{\min}(G)$ edges crossing it in G), then Poe thinks that C is a max-cut in graph G' (that is, C has $c_{\max}(G')$ edges crossing it in G'). Sadly this is not true; that is, C may not be a max-cut in G' . Find a small counterexample (G , which should be a simple graph) with at most 5 vertices. Where does Poe's logic go wrong? (It's sufficient to explain why this logic fails in your example.)

- (c) (7 points) In class we learned Karger's randomized min-cut algorithm as well as the AmplifiedKarger algorithm, which makes independent runs of Karger's algorithm on n vertices. Here we consider another variation of Karger's algorithm, which we call (a, b) -Karger and we will compare it to Karger's algorithm. (a, b) -Karger first contracts random edges until the multigraph contains a supernodes (the initial graph has n vertices), as Karger's algorithm normally does (except Karger's continues until 2 supernodes remain

and here we stop earlier). Then the algorithm makes b independent runs of Karger's algorithm on this multigraph on a supernodes. The algorithm outputs the smallest cut found in any of the b runs.

- (i) (1 point) Suppose we run Karger's algorithm twice independently and output the minimum cut found in either iteration. What is the probability of success; that is, with what probability does this algorithm find a min-cut in at least one of those two runs? (As we did in class, provide a lower-bound on the probability of success of this algorithm.)
- (ii) (3 point) Determine the total number of edge contractions performed by (a, b) -Karger (in terms of n, a, b). Give a lower bound on the probability of finding a minimum cut (in terms of n, a, b). (Assume $a < n$ for simplicity. Also, as we did in analysis of Karger's algorithm, you should lower-bound the probability of finding a specific min-cut for simplicity.)
- (iii) (3 points) Let us compare Karger's algorithm and (a, b) -Karger. Running Karger's algorithm twice as in part (i) above, results in $2(n - 2)$ edge contractions over both runs. In this problem, we will find values of a, b such that (a, b) -Karger performs $2(n - 2)$ edge contractions in total, but its probability of finding a min-cut is better than running Karger's algorithm twice independently. Fix $a = n^{1/3}$ (it turns out that this setting is optimal). For this setting of a , find the value β of b which leads to $2(n - 2)$ edge contractions. Then, using your answer from (ii), give a lower bound on the probability of success of the $(n^{1/3}, \beta)$ -Karger algorithm (use Ω to simplify your answer). (Hint: (a, b) -Karger's probability of success should be better than that of Karger's run twice in Part (i). You may find it useful to simplify your answer to Part (ii) by replacing $(a - 1)$ terms by a and $(n - 1)$ terms by n , if any.)

2 (10 points) Minimum Spanning Trees

In all of the following assume that G is an undirected, connected graph on n nodes with distinct, positive edge weights. Also assume that G is simple (no self-loops and no multi-edges) in all parts.

- (a) (2 points) Consider the edges of G sorted in increasing order of weight and place the first $n - 1$ edges into a set S . That is, S is the set of $n - 1$ edges with the smallest weights.

Prove or disprove: S is a minimum spanning tree in G (assume $n \geq 2$).

(Concisely prove the claim or give a small counterexample.)

- (b) (3 points) Prove or disprove: There is a unique minimum spanning tree in G . (Concisely prove the claim or give a small counterexample.)

- (c) (5 points) Let T be a minimum spanning tree of G . Let e be some edge in G (which may or may not belong to T). Let us obtain a new graph G' from G by preserving everything the same except that we decrease the weight of e (but assume that the weights of all edges in G' are still distinct). Design an algorithm that can find a minimum spanning tree T' of G' in time $O(n)$, given G, T, e and its newly decreased weight w . Prove the correctness and analyze the runtime.

3 (12 points) Barr Selling Fish

Barr the Bear has started a business to sell fish to Poe and his fellow penguins. The penguin customers submit many fish orders, but Barr can only process one order at a time. Suppose that Barr currently has orders from n penguin customers (label them as $1, 2, \dots, n$). Customer i 's order takes t_i time to complete. Barr is going to process each penguin's order one by one, and the scheduling of orders can be described as a permutation of the customers. Let C_i denote the completion time of order i . For instance, if customer j 's order is the first to be completed, then we would have $C_j = t_j$ (assume Barr begins processing orders at time 0); if customer k 's order is the second to be completed after that, then $C_k = C_j + t_k = t_j + t_k$, and so on. Each customer is of different importance to Barr's business, and we denote this relative weight by w_i (for customer i). Barr wishes to minimize the weighted sum of the completion times of n orders, $\sum_{i=1}^n w_i \cdot C_i$. Intuitively, the more important a customer is, the sooner Barr wishes to complete the customer's order.

Design an $O(n \log n)$ algorithm to solve this problem to help Barr. You are given a set of n orders with a processing time t_i and a weight w_i for each customer i (assume t_i, w_i are positive integers). You want to decide an ordering of the customer orders so as to minimize the weighted sum of the completion times. Prove correctness of your algorithm (i.e., optimality) and analyze its runtime.

4 (12 points) Attack on Penguin

Do you know that there is a Bear Island in Antarctica? Barr the Bear has landed on Bear Island and is looking to invade Poe the Penguin's homeland. He has created bases in n islands, including Bear Island (we will label islands from 1 to n , where Bear Island is labeled as 1). Barr has also constructed m bidirectional transit lines between the islands for his covert operations. The line connecting island i and island j takes w_{ij} minutes to traverse and (coincidentally) costs w_{ij} Barr coins per hour to maintain. For each of the transit lines, $w_{ij} > 0$. Let d_i be the shortest time it takes to traverse from Bear Island to island i .

Barr the Bear spent too much money on honey and is running low on coins. He decides that he wants to close down some transit lines to save coins. However, he still wants to ensure that closing down transit lines will not affect the shortest amount of time to travel from Bear Island to each island (that is, even after closing down transit lines, d_i would still be the same).

Given n islands, the list of m transit lines, and their costs, help Barr the Bear find a subset of transit lines such that the cost of maintaining the lines is minimized (the total number of coins to be spent per hour), and yet the time taken to travel from Barr Island to every other island remains the same as before closing down any lines.

Solve the problem in $O(m+n \log n)$ time. Your algorithm should output which transit lines are remaining (or which ones to be removed, equivalently). Prove correctness and analyze runtime. If you are using a graph formulation, remember to explicitly state what is the graph you are using (e.g., what are the vertices and edges).

5 (0 points) Routing Packets of Fish

Note: This problem will not be graded (please do not submit as we will not read), but it is given as an exercise problem (since we will not cover Network Flow before this homework assignment is released).

Poe and Barr collected so many fish that they have more than they can eat! Therefore, they've decided to ship their fish to other penguin and bear communities across the poles. The n communities form a connected graph connected by a series of m ice chutes. Fish packets can travel across these ice chutes very quickly, but the shipping is limited by the chute's capacity. If communities i and j are connected by a chute, then we can pass $c(i, j)$ fish at once. We say that given a path of chutes, the capacity of the path is the minimum capacity over all the chutes along the path. Poe and Barr want to find paths through the network of chutes that maximize the capacity of fish they can ship.

- (a) (0 points) Suppose Poe and Barr set up their shipping depot at community k . Give an efficient algorithm to find for every community c , maximum capacity of any path from k to c . Prove your algorithm's correctness and runtime.
- (b) (0 points) Now suppose Poe and Barr want to facilitate trade amongst all of the polar communities. Give an $O(n^2)$ time algorithm to find for every pair of communities, maximum capacity of any path between them. You may assume that chutes are bidirectional with the same capacity in each direction. Prove your algorithm's correctness and runtime.

Bonus (0 points): Suppose Poe is at community p and Barr is at community b and they only care about shipping fish between one another. Give an $O(m)$ time algorithm for finding the path of maximum capacity between p and b .