

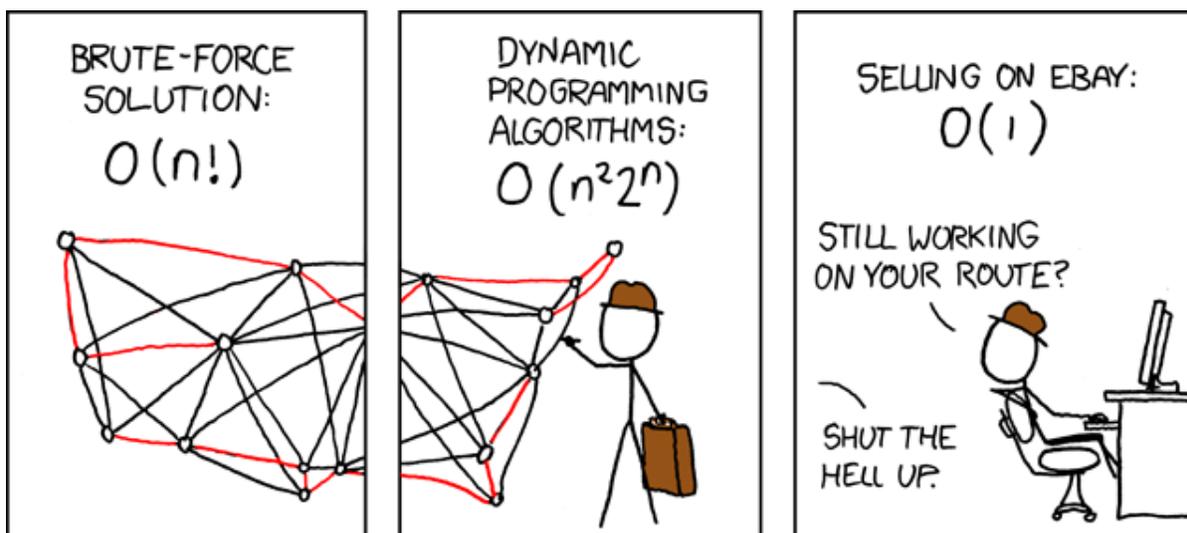
Dijkstra's Shortest Path Packet

Math-270: Fall of 2017

Prof. Gregory V. Bard

Contents:

Page 1	This cover page.
Page 2	Pseudocode for Dijkstra's Algorithm.
Page 3	A brief FAQ about deliverables (What is the answer? Is it Unique?)
Pages 4, 5, ..., 12	Example Graphs 1, 2, 3, 4, 5, 6, 7, 8, and 9.
Page 13	Some Information about a non-Navigational Example
Page 14	Example 5A (Air Freight)
Page 15	A spare copy of Example 1, for practice at home.
Pages 16, 17, 18	The main solutions: a list of the predecessor arrays, the distance arrays and the tables of shortest paths.
Online	The trees of shortest paths presented 2 ways.
Coming Soon	Proof that this actually works.



Dijkstra's Algorithm for the Shortest Paths in a Graph (or Digraph)

Math-270: Discrete Mathematics

April 18, 2017

Inputs: The set of vertices, V . The weights of the edges, as a two-dimensional array, where $\mathbf{weight}[i, j]$ is the weight of the edge from v_i to v_j . If no such edge exists, then denote this by $\mathbf{weight}[i, j] = \infty$. Also, some $v_a \in V$ is identified as “the headquarters vertex.”

Note: $\#V$, the number of vertices is usually passed as well, though this might be obvious from the size of the set V and the dimensions of the array \mathbf{weight} .

Outputs: The array $\mathbf{Distance}$, which has the distance from each vertex to v_a . Note that $\mathbf{Distance}[v_a] = 0$, and any vertices not reachable from v_a have distance ∞ . Also, the array \mathbf{Pred} , a list from which the shortest paths from v_a to any other vertex can be easily computed.

1. $\mathbf{Distance}[v_a] \leftarrow 0$.
2. For each $v \in V - \{v_a\}$ do
 $\mathbf{Distance}[v] \leftarrow \infty$.
3. $Q \leftarrow V$.
4. $D \leftarrow \{\}$.
5. While $Q \neq \{\}$ do
 - (a) Select the $u \in Q$ such that $\mathbf{Distance}[u]$ is the smallest.

Note: In the first iteration, this is always v_a .

- (b) (rare) if $\mathbf{Distance}[u] = \infty$ then the graph is disconnected; Q is the connected component containing v_a . Viewing Q as a subgraph, you've solved the problem on Q , and you can't reach any of the other vertices at all. Either report an error message or throw an exception.

- (c) $D \leftarrow D \cup \{u\}$.

- (d) $Q \leftarrow Q - \{u\}$.

- (e) For each $v_i \in \text{neighborhood}(u)$ do

- i. $\mathbf{CandidateDistance} \leftarrow \mathbf{Distance}[u] + \mathbf{Weight}[u, v_i]$. *******

- ii. If $\mathbf{Distance}[v_i] > \mathbf{CandidateDistance}$ then

- $\mathbf{Distance}[v_i] \leftarrow \mathbf{CandidateDistance}$.
- $\mathbf{Pred}[v_i] \leftarrow u$.

Note: At this point, $Q = \{\}$.

6. Return the arrays $\mathbf{Distance}$ and \mathbf{Pred} .

Q&A about “What’s the Answer” and “Is the Answer Unique?”

Q: When this algorithm is run, what is the answer?

A: This is a real-life algorithm, so there are several outputs. “Which one is the answer” depends on the situation/application. This is not an algebra problem where you write down $x=4$ at the end of some equation-solving.

- **The array of distances** – This is certainly useful information, but usually you want something more than this. You cannot navigate by this, for example.
- **The predecessor array** – In some ways, this is *only* useful because it is used (in the back-tracking sub-algorithm) to generate the list of shortest paths. However, it has other advantages: in many computer languages, it is much easier to pass this simple array around rather than a much more complicated data structure. It is also much easier to grade this than the list of shortest paths, so some instructors might prefer it.
- **The list of shortest paths to each vertex** – This is the primary output. Using this, you can navigate from the origin vertex to any other vertex, confident that you are taking the shortest path to your destination.
- **The tree formed by the union of all the shortest paths** – If you take all the edges used in the shortest paths from one common origin to all possible destinations, then you get a tree. For some navigation problems, this tree would be very useful information. Ant colonies know about this aspect of shortest-path calculations, and that is why you will sometimes see “ant highways” when an ant colony has found abundant food nearby.

Q: How are the answers shown in this packet?

A: In this packet, the first set of answers will provide the predecessor array and the array of distances, as well as the shortest paths to each vertex. Next, all the spanning trees are presented by highlighting the relevant edges in light green, on the original graphs. Last, all the spanning trees are presented again, by drawing them to look like trees.

Q: How should I write the answer on an exam, a test, a quiz, or a homework assignment?

A: You have to actually read the instructions. For example, if you are asked to write the predecessor array, then you should write the predecessor array. If you are asked to make a list of shortest paths, then you should make a list of shortest paths. For this reason, make sure you know the names of all the different answer formats.

Q: Does Dijkstra’s Algorithm compute the minimum spanning tree?

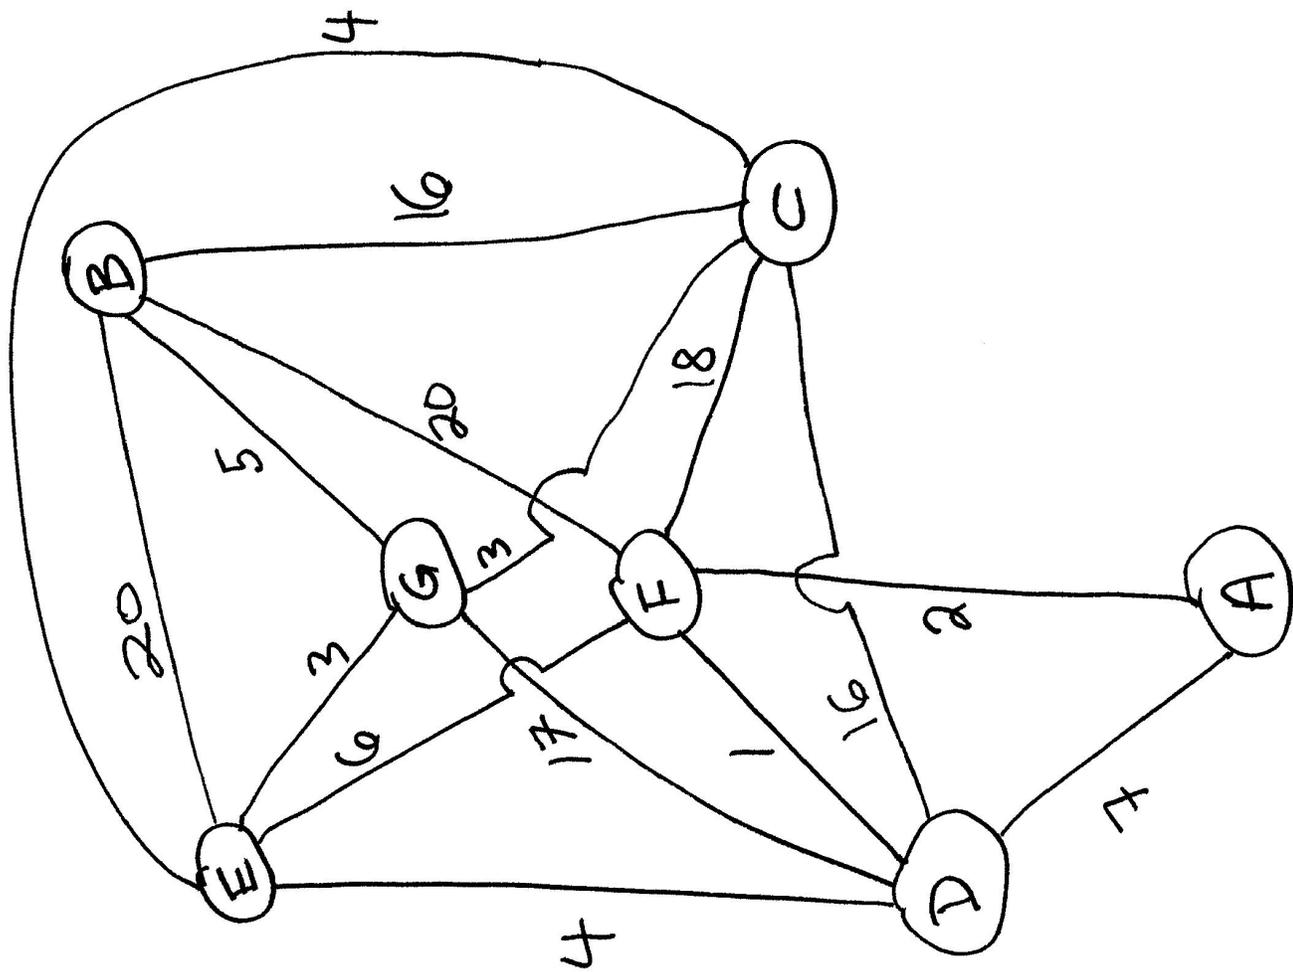
A: No. The tree of shortest paths is a spanning tree. That means it is a tree (a connected acyclic graph) that includes all the original vertices, and is comprised only of edges that exist in the original graph. It will have the minimum weight of all possible spanning trees *that have the origin vertex as the root of the tree*. However, there could be some other spanning tree, with a different root, that has a lower minimum total weight. When we ask for a minimum spanning tree, we want to consider all possible roots. That is a question for which Dijkstra’s algorithm is not suitable.

Q: Are the answers unique?

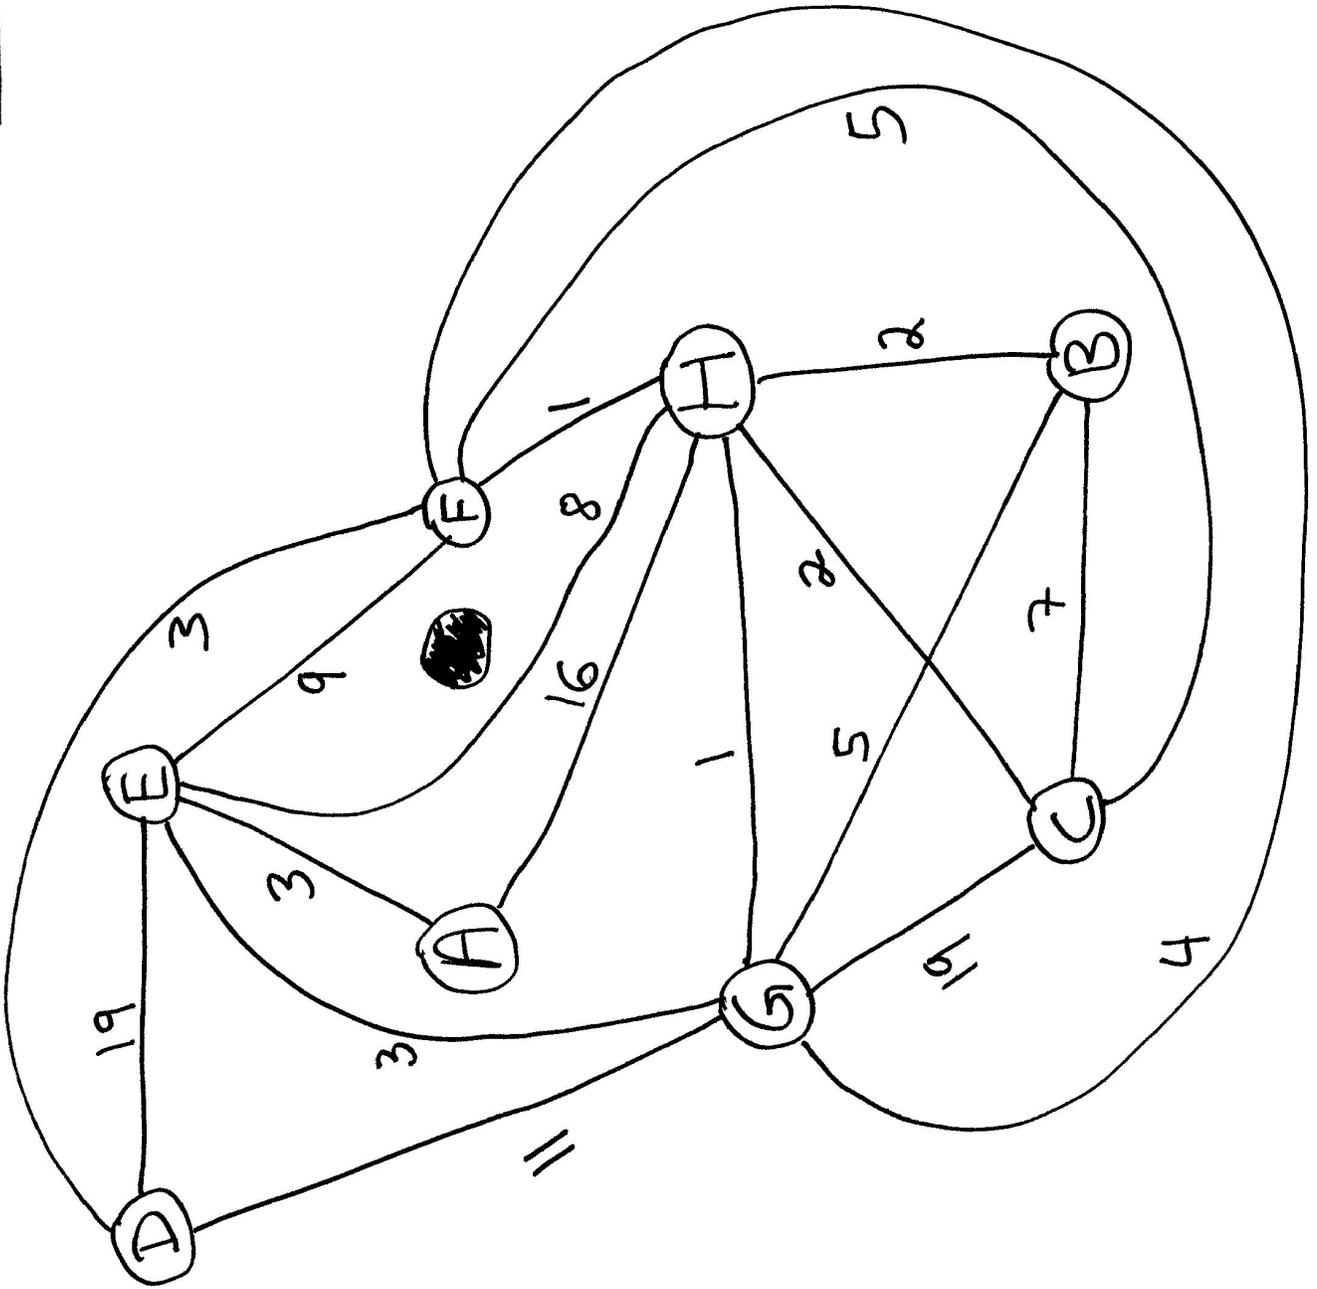
A: Questions of uniqueness (or existence & uniqueness) are favorites of pure mathematicians. As it turns out, if there are no ties in the algorithm, then all the answers are unique. By a tie, I mean the step where we “select the u from Q for which $\text{distance}[u]$ is the smallest.” If this selection never has a tie for smallest, then all the answers are unique.

Even if there is a tie, then the array of shortest distances will always remain unchanged. If you and your classmate carry out the algorithm and break the tie in different ways, then you will possibly get different paths to some vertices, but the distances along those paths are the same. The spanning tree might also be different, but it remains a tree (a connected acyclic graph) and it has the minimum weight of all possible spanning trees whose root is the vertex that you were starting the algorithm from.

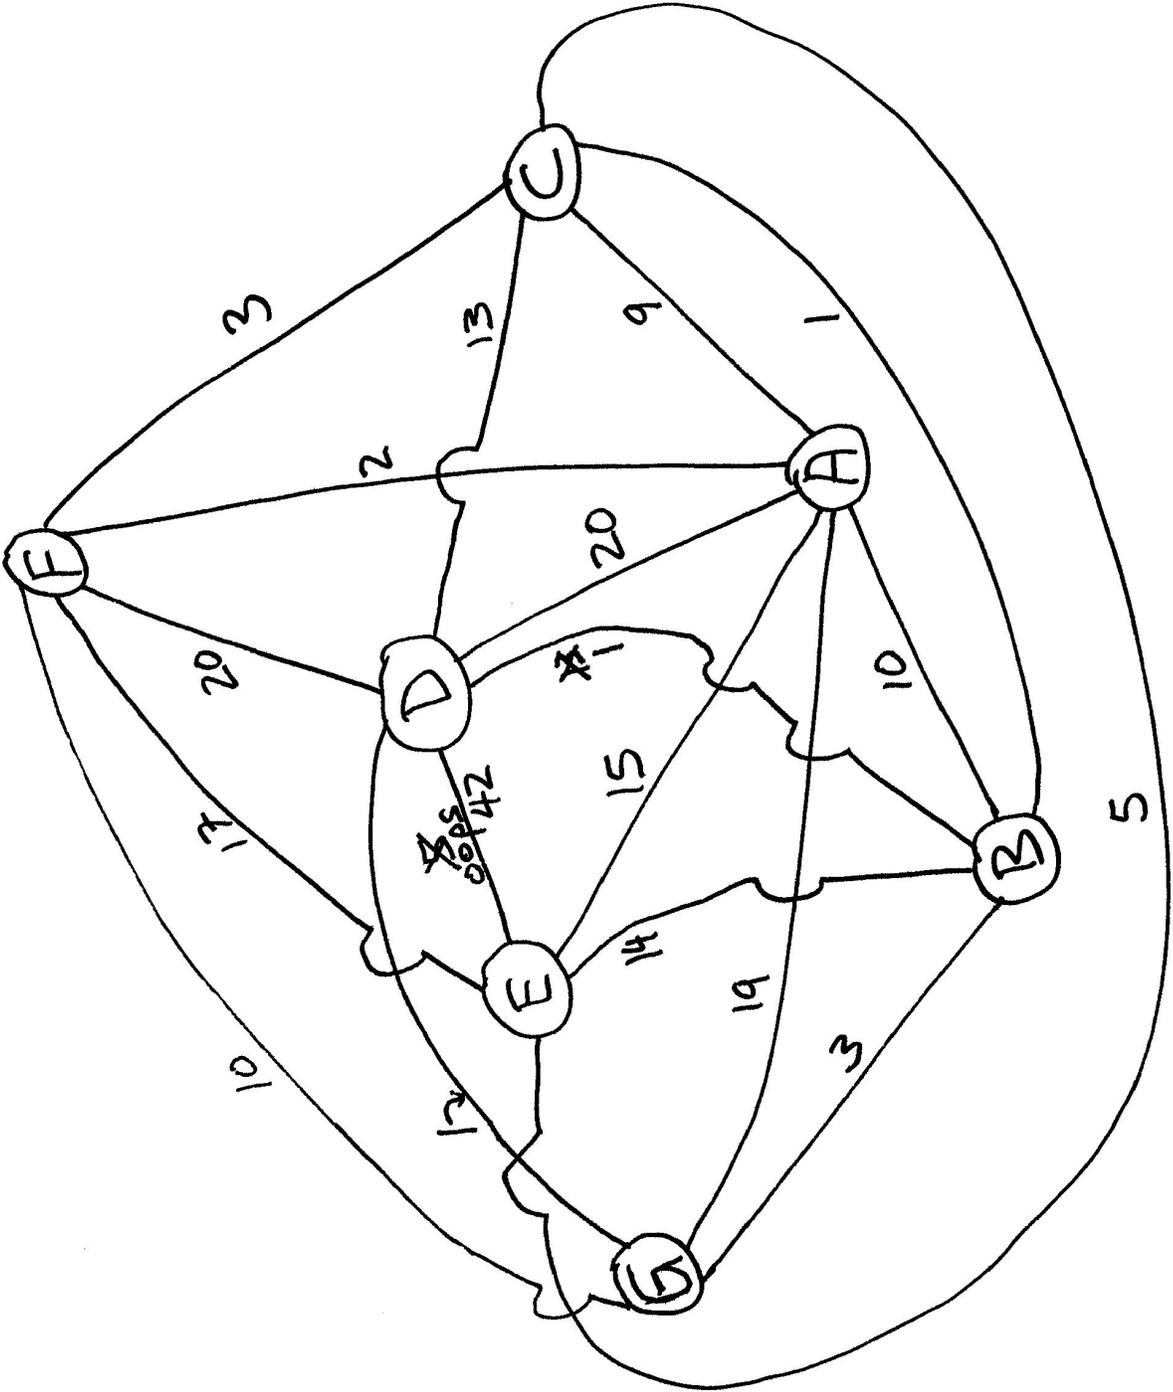
Ex 1:



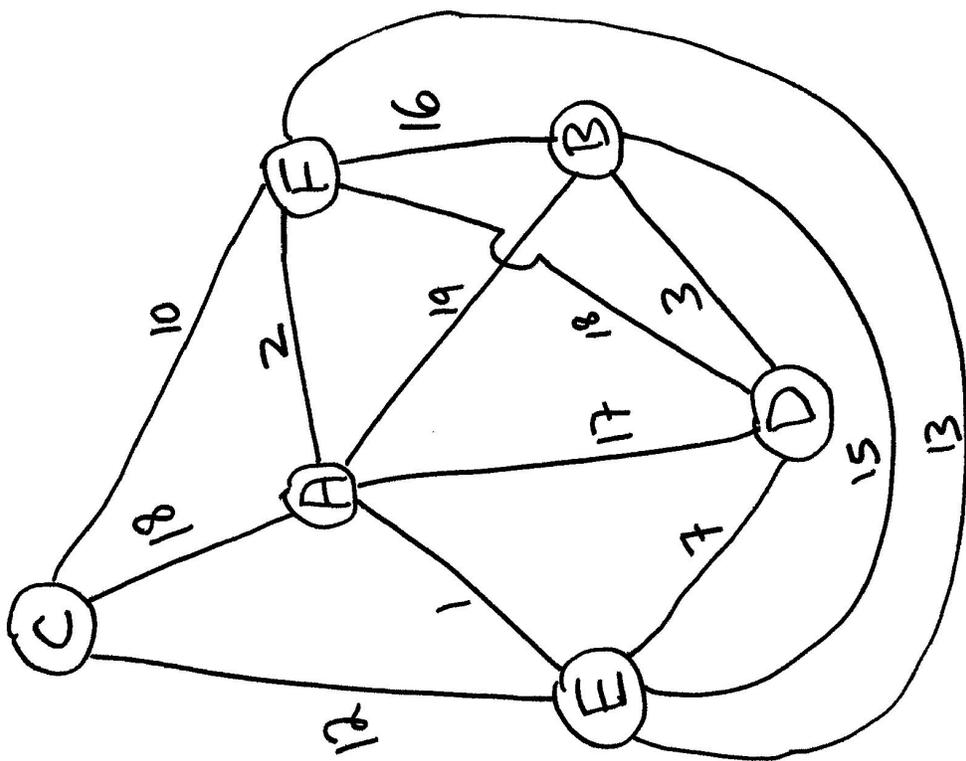
Ex 2



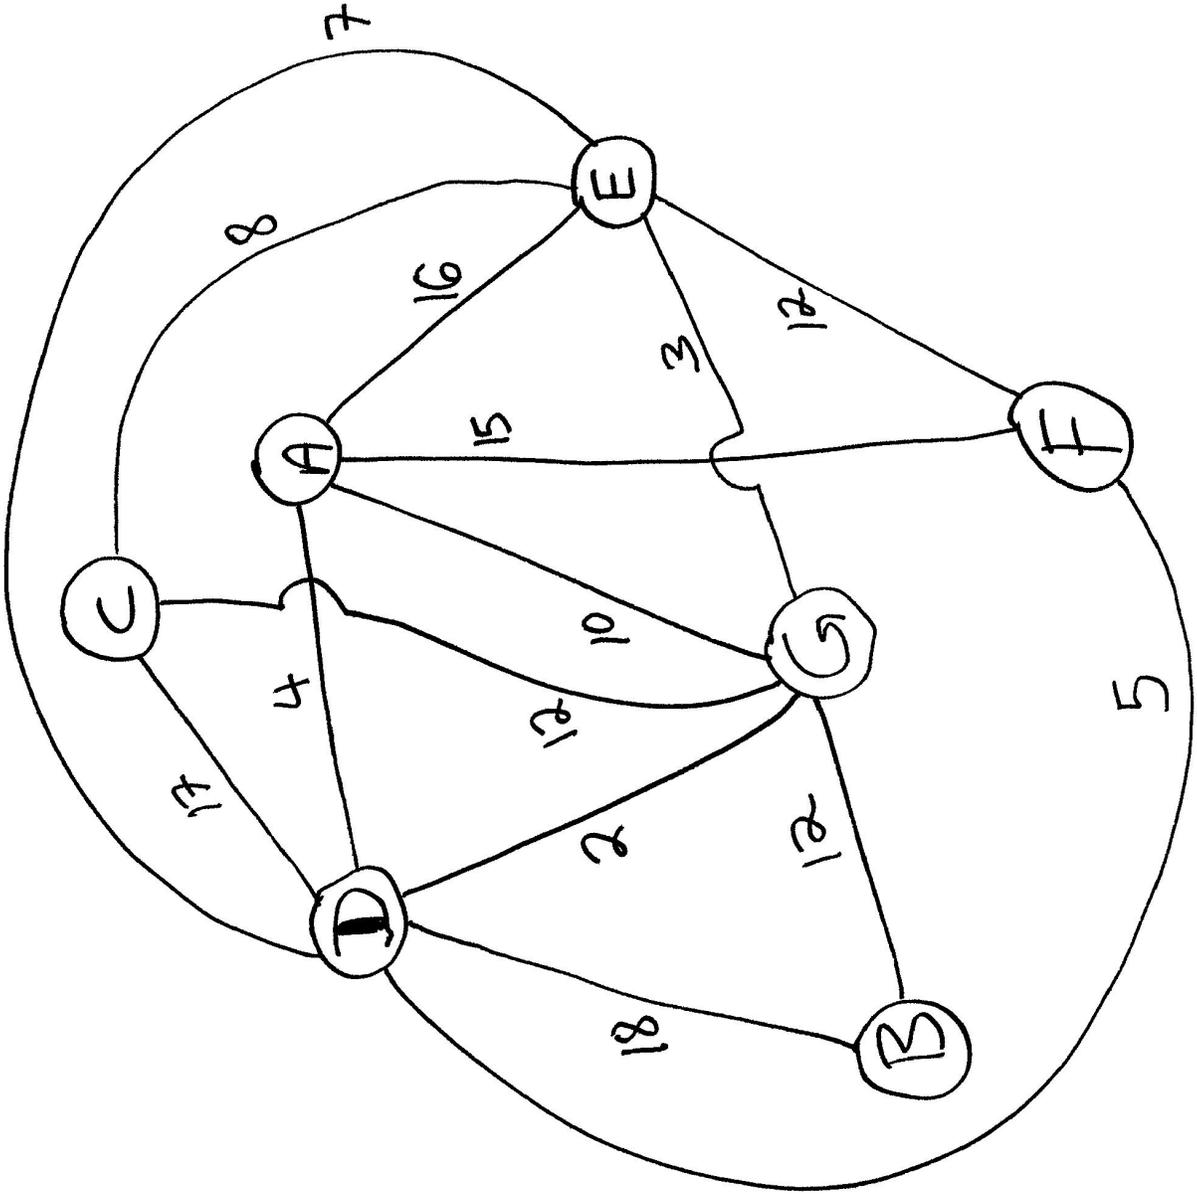
Ex 3



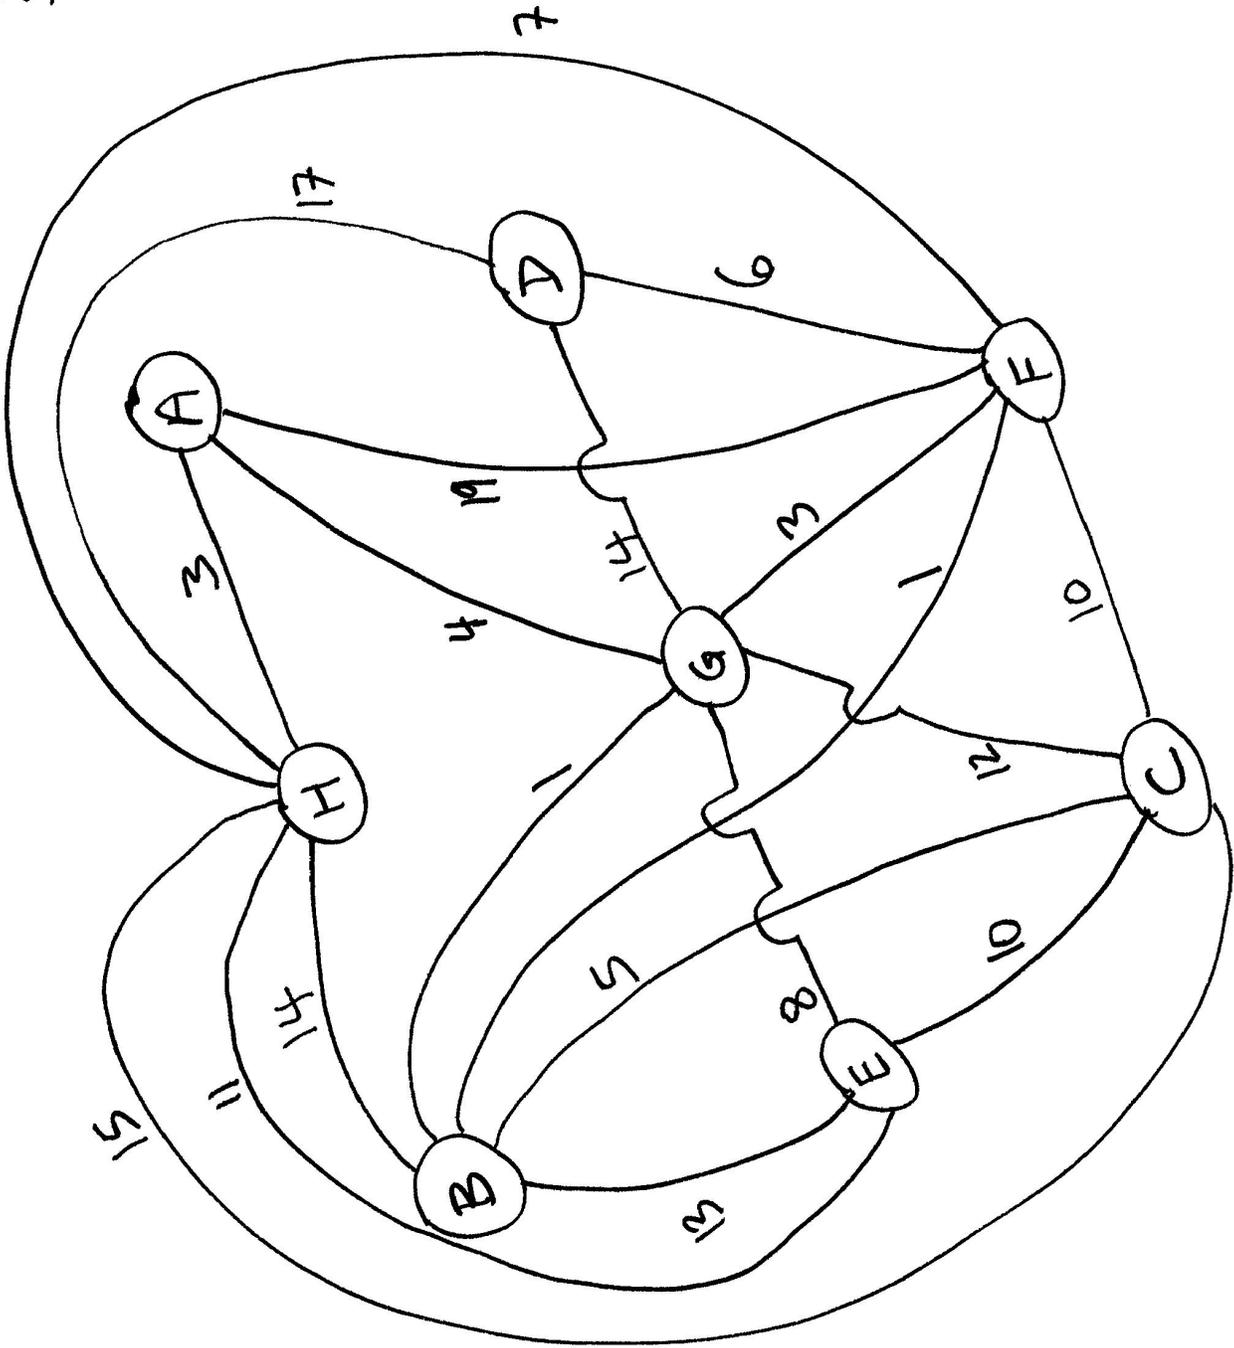
Ex 4



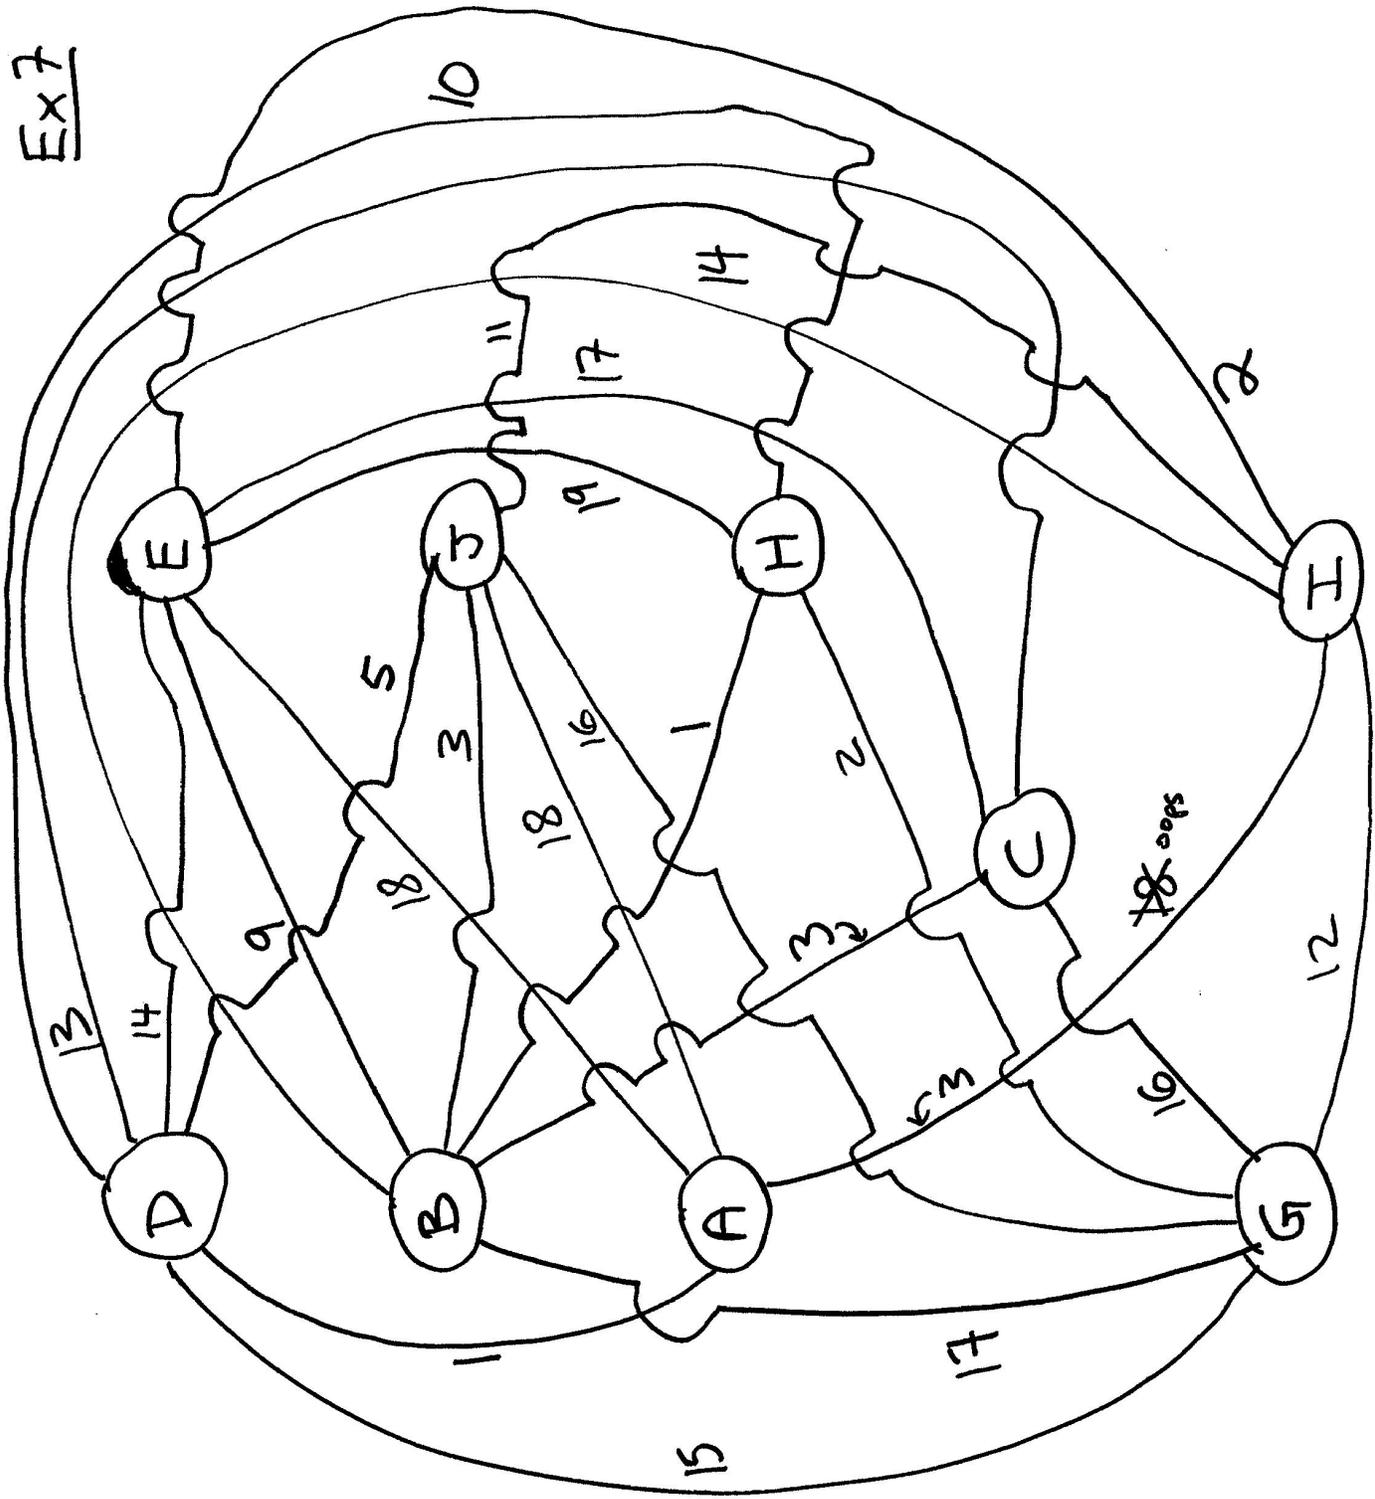
Ex 5



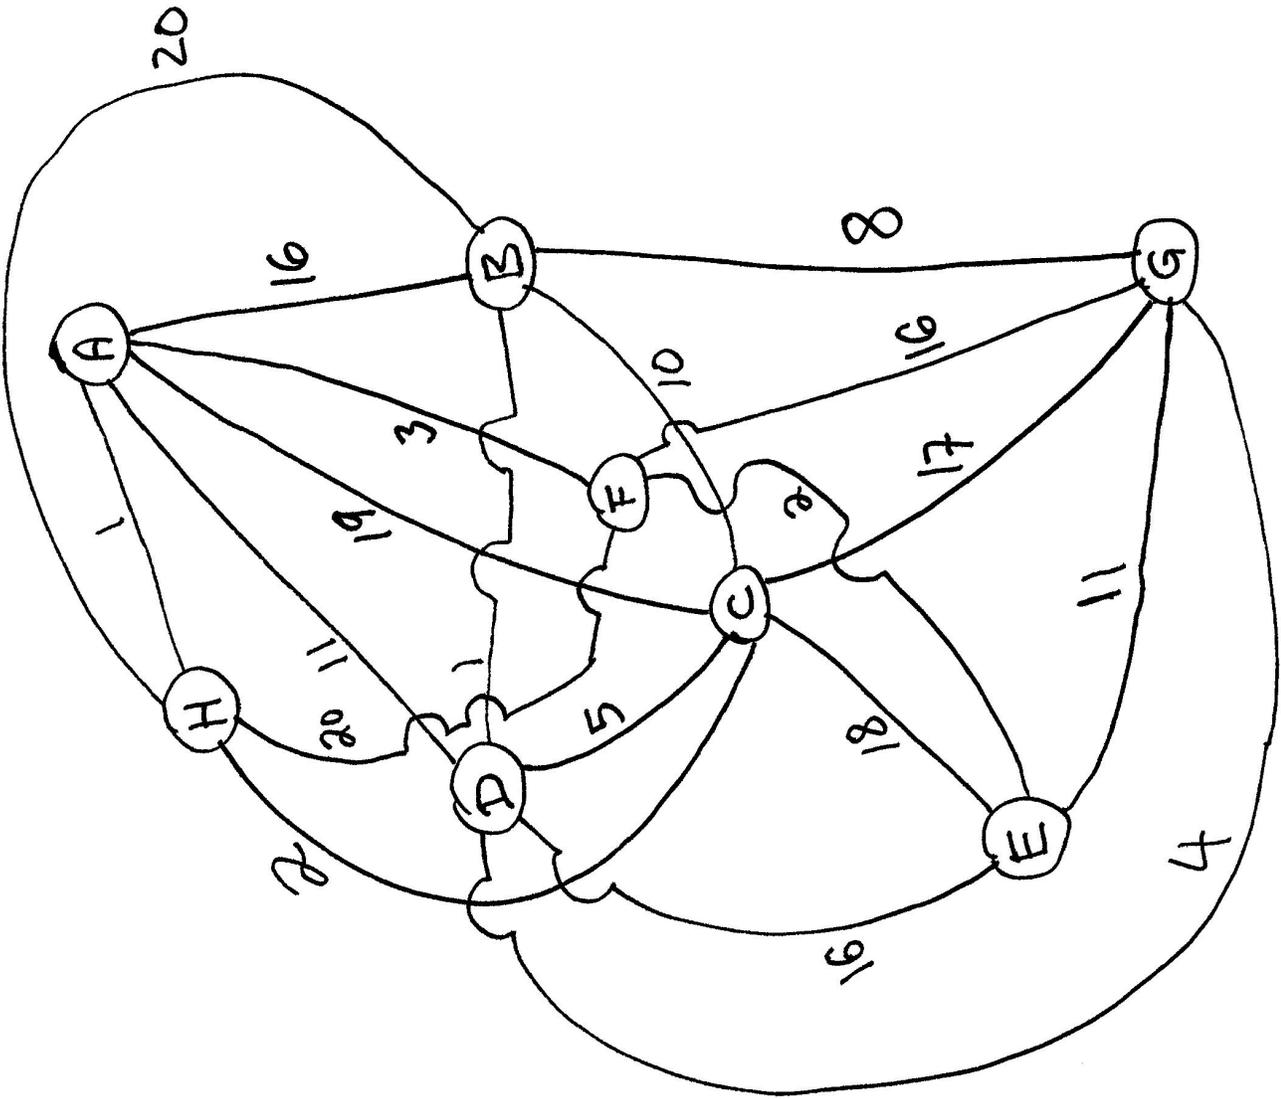
Ex6



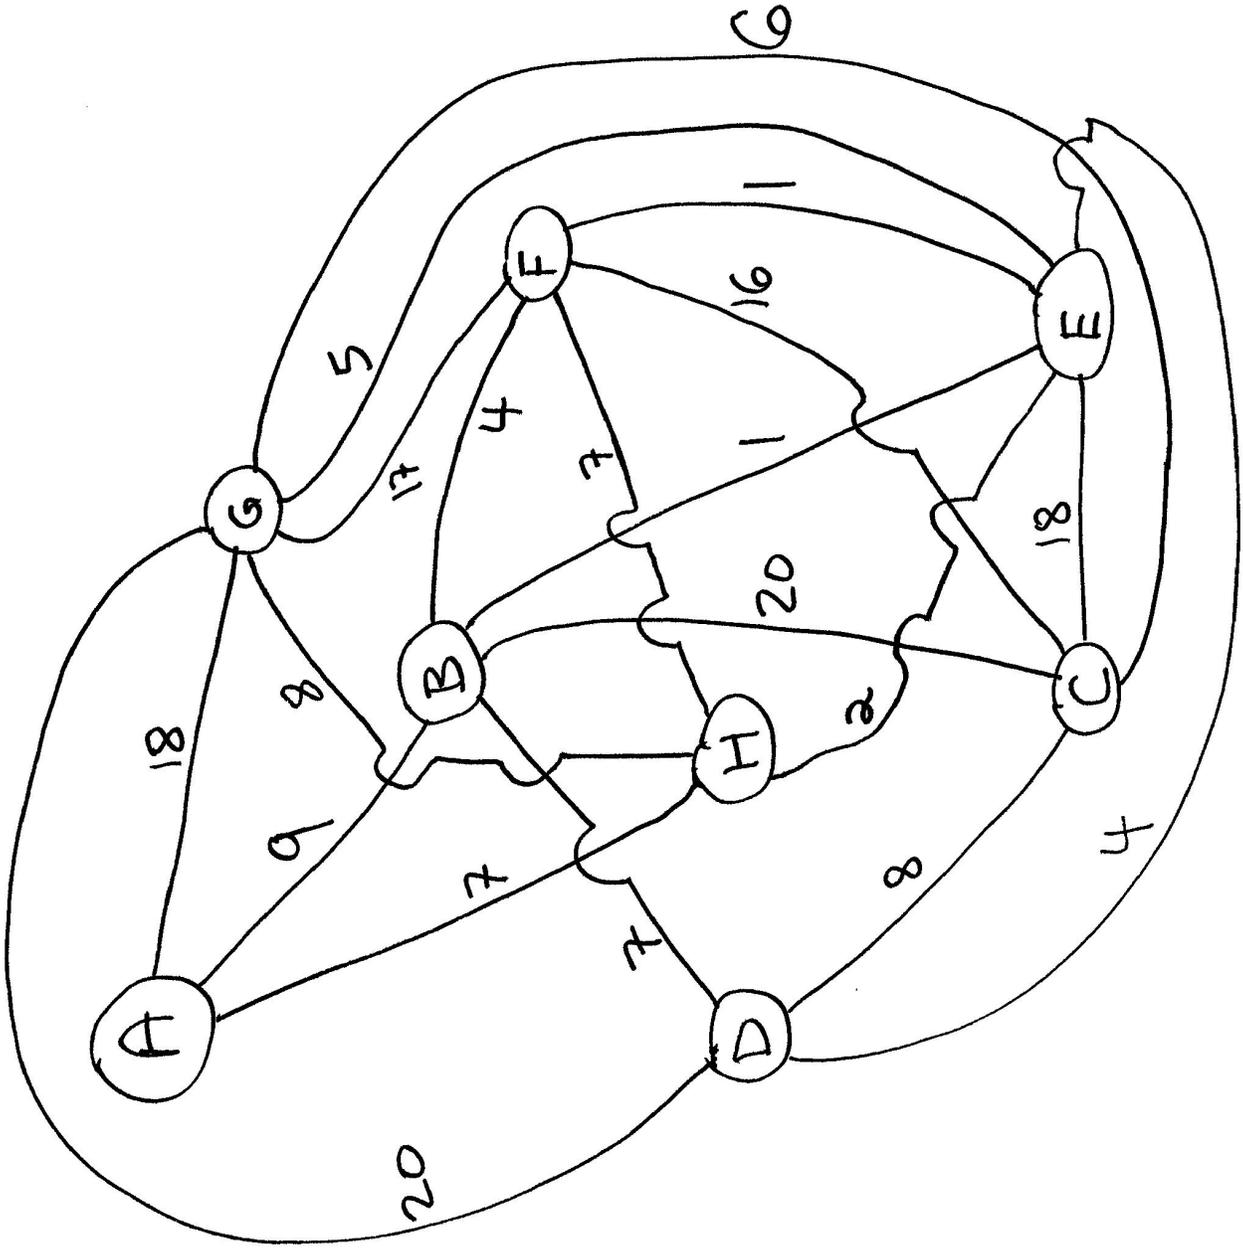
Ex 7



Ex 8



Ex 9



A Non-Navigational Application

After Example 9, you will see Example 5A, and you might be wondering what that's about. I decided that it would be nice to have an example that is not about navigating and finding the shortest distance.

This problem represents an air-freight problem. The company, headquartered in Atlanta, has customers throughout the upper Midwest, and also in Dallas. Because airfares change suddenly and without notice, their software would have to run the Dijkstra Shortest Path algorithm extremely often---every time a fare changes.

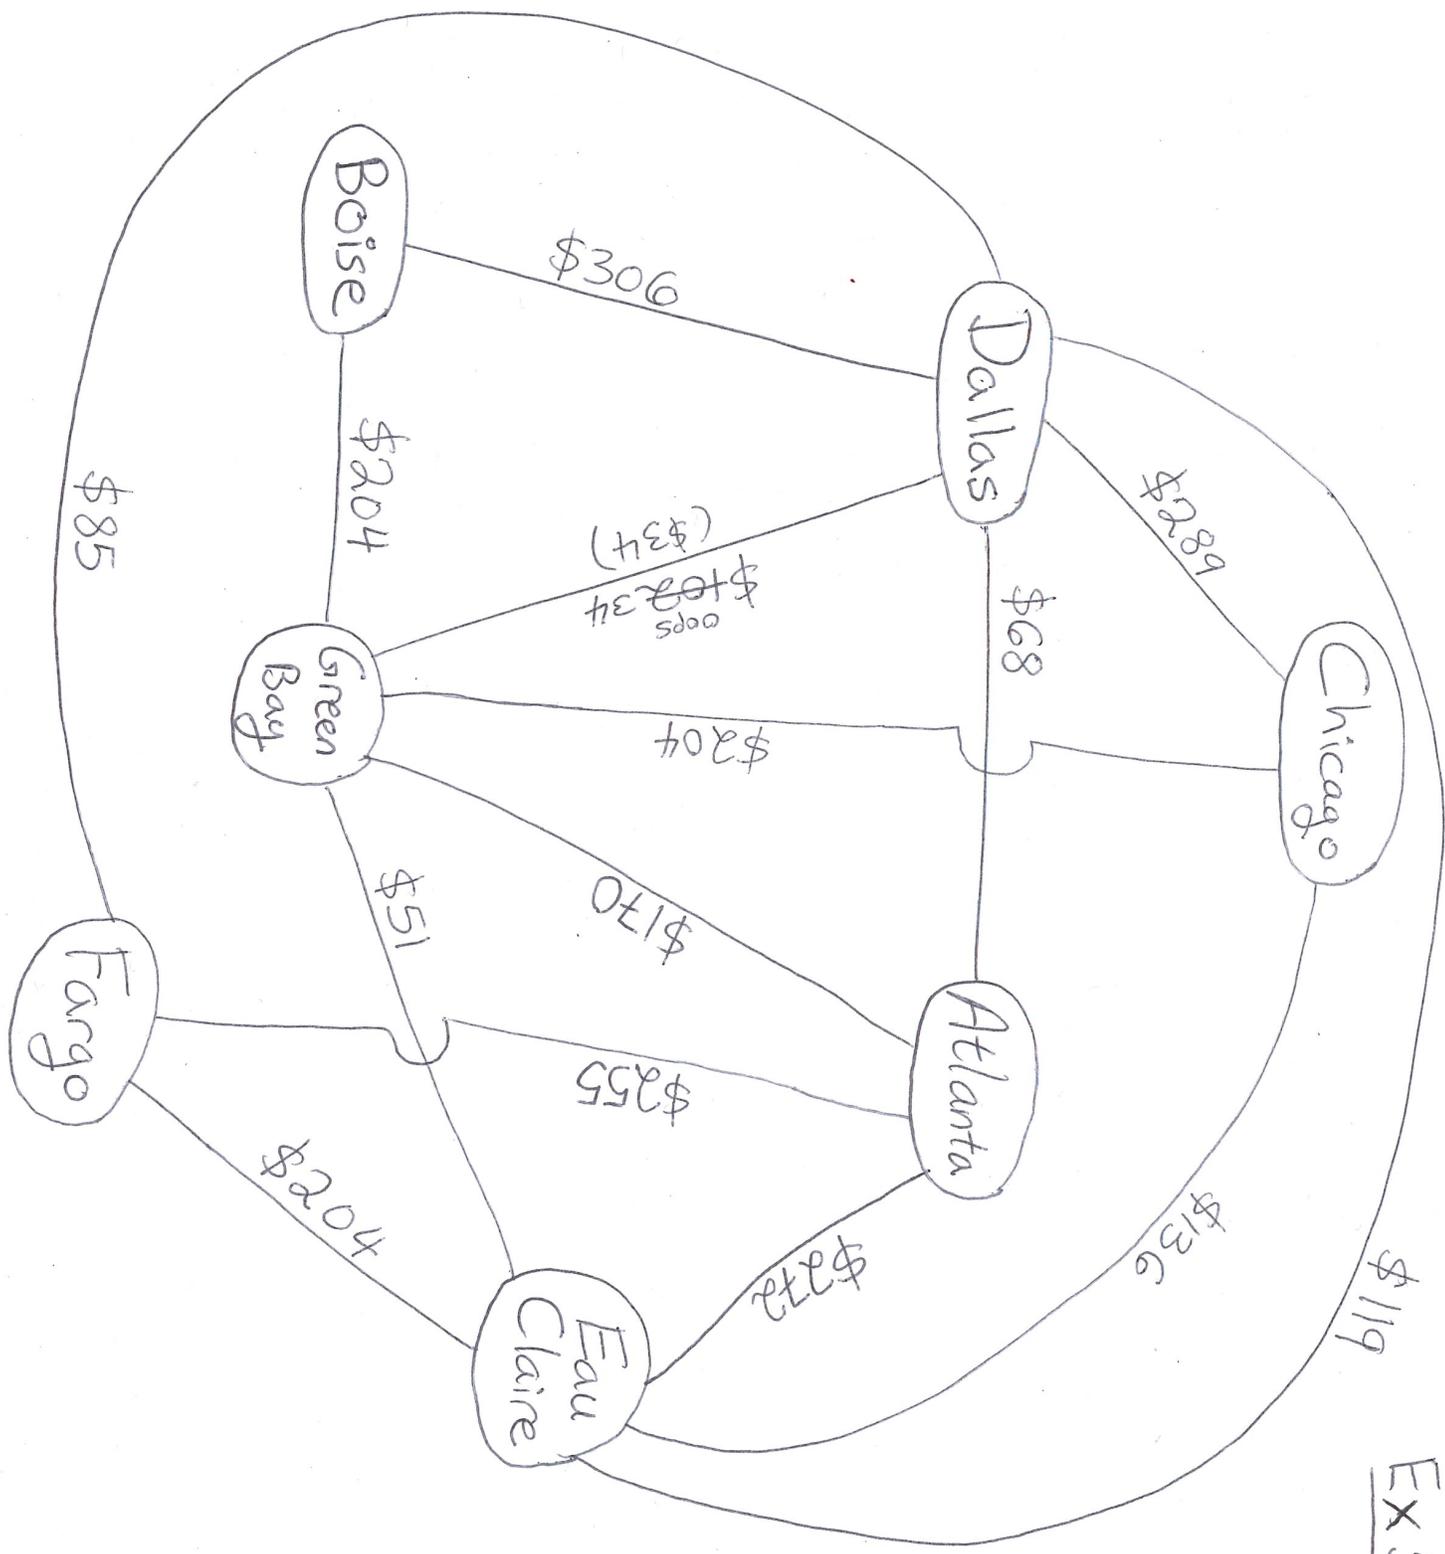
In this problem, an extremely "good deal" has come up. The fare from Dallas to Green Bay is far cheaper than expected. It is roughly 1/9 the cost of Dallas to Boise, even though the distance is roughly comparable. When you run the algorithm, you will see that this will transform Green Bay into a virtual hub for the company. All packages going to Boise, Chicago, Eau Claire, or Green Bay will be sent to or through Green Bay, at least as long as this bargain lasts. (That's 4 out of 6 possible destinations!)

Humans are not so good at realizing all the ramifications of one price change in a network of prices. That's the advantage of coding an algorithm, like Dijkstra's algorithm, that is provably correct.

It should be noted that it is important that this is air freight and not the flight of human passengers. Usually a human would prefer a flight with 0 or 1 stops to a flight with 3 or 4 stops. If I have a possible itinerary where I have to change planes 3 times, but it is \$200 cheaper than a direct flight, for sure I will choose the direct flight. That preference is hard to model mathematically, but with air freight, we would just want to know the cheapest way of sending a box of cargo from one point to another. The box of cargo does not get the chance to provide us with preferences about its flight path.

Last but not least, you might be wondering how I came up with the graph and the weights. As it turns out, I just took Example 5, and renamed the vertices. The weights are simply the weights from Example 5, but multiplied by \$17 to make them look like realistic costs. This was done to emphasize that Dijkstra's algorithm solves this problem without modification. You could code an air-freight routing system for this company using what you are learning about Dijkstra's algorithm, except that a real company would have 50-150 airports that it is servicing.

EX 5A



Ex.1:

