

Assignment 7: Pathfinder

Assignment by Julie Zelenski and Eric Roberts

Have you ever wondered how MapQuest and Google Maps or those tiny little GPS units work their magic? Get ready to find out. In this final CS106B challenge, you will create the Pathfinder application, which (in addition to other useful functions) finds and displays the shortest path between two locations on a map. The Pathfinder assignment is designed to accomplish the following objectives:

- To give you practice working with graphs and graph algorithms.
- To offer you the chance to design and write a class of your own.
- To let you practice using callback functions in an application.

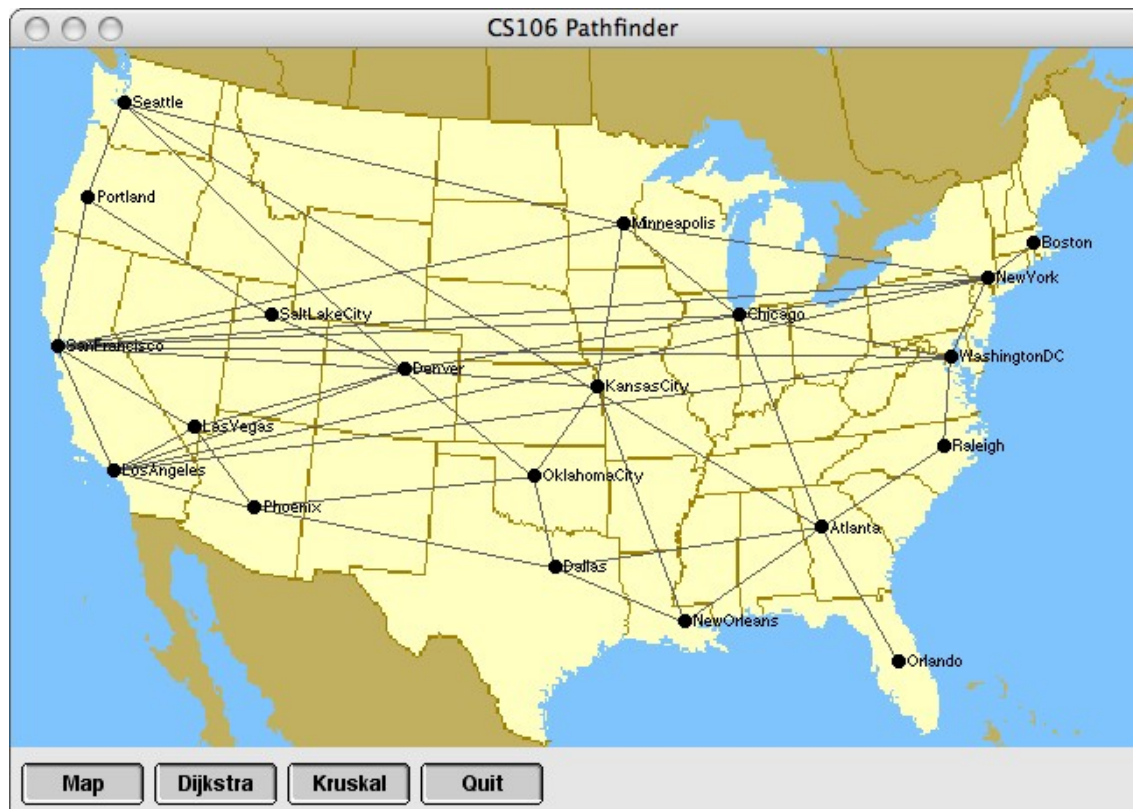
**Due Tuesday, June 12 at 11:30 AM.
No Late Submissions will be Accepted.**

You cannot use late days on this assignment. This assignment is due during the time in which we'd normally have the final exam, and university policy prevents us from having assignments due at any point after that.

The Pathfinder application from the user's perspective

The Pathfinder application begins by asking the user to specify a filename containing data for a map consisting of a background image and a graph linking locations in the map. The application then displays the graph in the graphics window. For example, if the user asks for the map `USA.txt`, Pathfinder will create the display shown in Figure 1.

Figure 1. Pathfinder after reading in the `USA.txt` data file



User interaction in the Pathfinder program is controlled by the buttons in the control strip at the bottom of the window. The **Map** button, for example, allows the user to enter the name of some other map, which it then loads into the window, replacing the existing one. The **Dijkstra** button executes Edsger Dijkstra’s algorithm for finding the shortest path between two nodes in the graph, as described in section 19.6 of the reader. For example, if the user selects **Dijkstra** and then clicks on the city circles for Portland and Atlanta, the Pathfinder application will highlight the shortest path between those cities, as shown in Figure 2. The **Kruskal** button executes an algorithm by Joseph Kruskal—described later in this handout—that finds the lowest-cost cycle-free subgraph connecting all the nodes. Clicking the **Quit** button exits from the program.

Figure 2. Result of applying Dijkstra’s algorithm to connect Portland and Atlanta



The program from an implementer’s perspective

The Pathfinder application is a substantial piece of code comprising seven files, if you count both headers and implementations. The good news is that you will need to change four of those files: **pathfinder.cpp**, **graphtypes.h**, **path.h**, and **path.cpp**. You will, however, need to use the rest of the files and must therefore understand their public interfaces, even if you can ignore the implementations for the most part. You also need to use some additional interfaces from the Stanford library, such as the **pqueue.h** interface that implements priority queues and the **GPoint** class from the **gtypes.h** interface. The complete list of files appears in Figure 3 on the next page, which offers an overview of how the individual pieces of the project fit together.

Figure 3. Files in the Pathfinder application

<code>pathfinder.cpp</code>	This file is the main program for the Pathfinder application. Most of your work in this assignment comes from adding code to this file; the version of <code>pathfinder.cpp</code> in the starter folder simply initializes the graphics library and puts the <code>Quit</code> button on the screen so that you have one example of how to use buttons. You have to write everything else. Your implementation of Dijkstra's and Kruskal's algorithms will appear here, along with the code for reading the graph data from the file and for responding to the other buttons.
<code>path.h</code> <code>path.cpp</code>	In the finished Pathfinder project, these files will define the interface and implementation of a <code>Path</code> class that represents a multistep path in a graph. As they appear in the starter project, however, these files are almost entirely empty, which means that you are responsible for both the design and the implementation of the class. This part of the assignment is described in the discussion of Phase 6 (which covers designing the <code>Path</code> class) on page 11.
<code>graphtypes.h</code>	This file is simply the structure-based graph interface that appears in Figure 19-3 on page 802 of the reader. You won't need to use the <code>SimpleGraph</code> type, because Pathfinder uses the object-based <code>Graph</code> class instead. You may need to add fields to the <code>Node</code> and <code>Arc</code> structure types to keep track of the additional data the Pathfinder application needs.
<code>gpathfinder.h</code> <code>gpathfinder.cpp</code> <code>gpathfinderimpl.cpp</code>	These files provide a library of graphical functions for the Pathfinder assignment. This interface includes all the calls that you need to make in implementing Pathfinder, so there is no reason—unless, of course, you are extending the assignment to add new features—that you would need to use the <code>graphics.h</code> or <code>extgraph.h</code> interfaces. As is typical in commercial projects, however, we have provided very little documentation of the Pathfinder graphics package in this handout. To find out what's available, you need to read the <code>gpathfinder.h</code> interface and use the comments to figure out what you need to know.

With Pathfinder—as with any large project—the most important advice we can offer is to subdivide the project into phases so that you can implement and test each phase independently. In general, if you write the code for Pathfinder all at once and then try to get it working, the complexity of the project will almost certainly make your life difficult as you debug your code. On the other hand, if you take the time to test each piece of your code as you write it, the assignment will go much more smoothly.

We recommend that you implement Pathfinder in the following phases:

1. Determine what extensions you need to make in the graph data structures.
2. Write the code to read map data from a file into its internal representation as a graph.
3. Use the facilities from `gpathfinder.h` to display the map on the screen.
4. Add buttons to the control strip, making sure that they call the appropriate functions.
5. Reimplement Dijkstra's algorithm so it fits into the Pathfinder application.
6. Design and implement the `Path` class and integrate it into the Dijkstra code.
7. Implement Kruskal's algorithm for finding minimum spanning trees.

The sections that follow describe each of these phases in turn.

Phase 1: Determine what extensions you need to make in the graph data structures

The Stanford C++ libraries include a `graph.h` interface that exports a parameterized `Graph` class that takes the node and arc types as template parameters. The `graphtypes.h` interface in the starter folder defines two structure types—`Node` and `Arc`—that define the minimal structures required to represent any graph. As supplied, these structures contain no information specific to Pathfinder. Depending on the design of your code—and also on what extensions you choose to implement—these types will need to include additional information. If nothing else, you won't be able to display the nodes on the graphics window unless you know the coordinates at which that node appears, which means that this information needs to be part of the data structure. You therefore need to change the `graphtypes.h` interface so that the structure types contain whatever information your application needs.

Beyond the additional data required for nodes and arcs, it is also likely that you will want to associate additional information with the graph as a whole. Because the `Graph` class is precompiled as part of the Stanford libraries, you can't change its definition. The good news, however, is that you can extend the behavior of `Graph` by defining a subclass that adds any additional data you need. Defining a `Graph` subclass has the further advantage of allowing you to incorporate the template parameters as part of the base type. If you don't define a new class, everywhere you need to use a graph, you have to include the type parameters, as in `Graph<Node, Arc>`. If, however, you create a new class with the header line

```
class PathfinderGraph : public Graph<Node, Arc>
```

you no longer need to specify the template parameters when you are working with the `PathfinderGraph` class.

Your mission in Phase 1 is to read through the assignment, figure out what additional information you need to maintain at each level of the graph structure, and then make the necessary extensions to `graphtypes.h` and your new `PathfinderGraph` subclass to incorporate that information into the data structure.

Phase 2: Read map data from a file into its internal representation

The next step is to write the code necessary to read the information defining a Pathfinder graph from a data file. The starter folder includes four data files: `USA.txt`, `Small.txt`, `Stanford.txt`, and `MiddleEarth.txt`. Figure 4, for example, shows the contents of the `Small.txt` data file, along with a few explanatory notes. With only four nodes and six arcs, this file is ideal for testing.

Figure 4. The contents of the `Small.txt` data file

<code>USA.jpg</code>	<i>Name of image file to display background picture</i>
<code>NODES</code>	<i>Marks the beginning of list of nodes</i>
<code>WashingtonDC 536 176</code>	<i>Each city is a one-word name with x-y coordinates</i>
<code>Minneapolis 349 100</code>	
<code>SanFrancisco 26 170</code>	
<code>Dallas 310 296</code>	
<code>ARCS</code>	<i>Marks beginning of list of arcs</i>
<code>Minneapolis SanFrancisco 1777</code>	<i>Each arc specifies two nodes and a distance</i>
<code>Minneapolis Dallas 935</code>	<i>Note that each arc is a bidirectional connection</i>
<code>Minneapolis WashingtonDC 1600</code>	
<code>SanFrancisco WashingtonDC 2200</code>	
<code>Dallas SanFrancisco 1540</code>	
<code>Dallas WashingtonDC 1319</code>	

Each of the data files shares a common data format. The first line is the name of a file containing the background image. The next line consists of the word **NODES**, which indicates the beginning of the node entries. The nodes are listed one per line, each with a name and the x and y coordinates of the node, with each of the fields separated by spaces. The line **ARCS** indicates the end of the node entries and beginning of the arc entries. Each arc identifies the two endpoints by name and gives the distance between the two. The end of the **ARCS** section is simply the end of the file.

Reading the Pathfinder data files is similar to reading the various other data files you've used this quarter. In general, the simplest approach is to read each line from the file using `getline` and then use a scanner to interpret the information on each line. As you go through the file, you also have to make the appropriate calls to `addNode` and `addArc` methods in the `Graph` class to create the appropriate data structure.

As you create the graph structure, one point to keep in mind is that the arcs in the data file are *bidirectional*, in the sense that each arc indicates a connection both from the first node to the second and from the second node back to the first. Because the `Graph` class itself assumes that arcs run in one direction, you will need to add two arcs for each line in the data file: one from the first node to the second and one in the opposite direction.

Given that Phase 2 comes before you have any means of displaying the graph on the screen, you won't be able to test it easily unless you write additional code to display the graph. That code is not required in the final assignment, but it is good practice to leave this kind of code in place so you can use it again if, for example, you need to change Pathfinder to work with a different file format.

Phase 3: Display the map on the screen

In Phase 3, your mission is to write whatever code is necessary to take a graph of the sort you've created in Phase 2 and display it on the screen. For example, assuming that you have read the graph data from the file `USA.txt`, there needs to be some function you can call that will produce the output shown back in Figure 1 on the first page of this handout. Such a function, however, is not sufficient in and of itself. At some point, you will need to be able to highlight individual nodes and arcs by drawing them in a different color, as shown in Figure 2. If you think about what drawing capabilities you are going to need and design this part of the assignment so that those capabilities are easy to achieve, you will have an easier time implementing the algorithms in the later phases of this assignment.

You are not, however, entirely on your own for Phase 3. We haven't made that much use of the `graphics.h` interface in CS 106B, and it wouldn't be fair to force you to figure out all the graphical facilities on your own. To make this assignment manageable, we've given you a `gpathfinder.h` interface that includes quite a few useful methods for implementing the graphical parts of this assignment. At least for the parts of this assignment before you start adding extensions, everything you need is available in `gpathfinder.h`, and you won't need to work with the `graphics.h` interface at all. In contrast to earlier assignments, however, we are not going to describe in detail what functions are available in `gpathfinder.h`. You need to look at the interface for that. In today's software development environment, programmers are *always* building on top of existing facilities whose structure they have to figure out on their own. To help you with that process, `gpathfinder.h` includes extensive comments, but you need to learn how to use it by looking at those comments and prototypes.

Phase 4: Add buttons to the control strip

For the last couple of decades, the Pathfinder application has used a text-based menu to determine what operation it should execute. That style of user interface dates from the 1970s and seems entirely out of place nearly forty years later. As you can see from the screenshots in Figures 1 and 2, the new Pathfinder has buttons, which at least drags it some distance into the modern world.

Although producing an application that seems a little less dated is useful in itself, the decision to incorporate buttons into this quarter's Pathfinder assignment has an additional advantage. In modern user interfaces, mouse clicks, button activation, and other similar user-generated activities represent actions that can occur at times not of the programmer's choosing. Such actions are collectively called *events*. Programs respond to those events by designating some function that should be called whenever that event occurs. The program as a whole typically does nothing except wait for events and then call the appropriate function.

Event-driven programs of this sort depend on the idea of function pointers (sometimes, as in Java, embedded in classes as dynamically bound methods but nonetheless implemented using function pointers) to trigger the action. The programmer begins by providing the event manager with a pointer to a function that must be called whenever a particular event occurs. At some later time, when the user performs an action that triggers the event, the event manager then uses that function pointer to trigger the appropriate action.

A simple form of this mechanism is illustrated by the contents of the `pathfinder.cpp` file in the starter folder. The main program in the starter file

```
int main() {
    initPathfinderGraphics();
    addButton("Quit", quitAction);
    pathfinderEventLoop();
    return 0;
}
```

where `quitAction` is defined as

```
void quitAction() {
    exit(0);
}
```

In this code, the call

```
addButton("Quit", quitAction);
```

creates a new button in the control strip and gives it the label **Quit**. More importantly, the call also registers the function `quitAction` as the function to be invoked whenever the user activates the **Quit** button. The `pathfinderEventLoop` function simply watches the mouse and waits for the user to click on one of the buttons. When the user selects a button, `pathfinderEventLoop` has to invoke the associated function. In the case of the **Quit** button, that function is `quitAction`, which invokes the standard `exit` function to terminate the application.

The `addButton` and `pathfinderEventLoop` functions are part of the `gpathfinder` module and have no knowledge of functions like `quitAction` that live in the main program. The implementations of `addButton` and `quitAction` lie on opposite sides of the abstraction barrier that the `gpathfinder.h` interface is designed to create. The function is provided by the main program and passed in pointer form across the abstraction barrier as a parameter to the `addButton` function. At some later time, the implementation of `pathfinderEventLoop` has to retrieve the function pointer and make a call back

across the abstraction barrier to the function defined in the main program. As noted in Chapter 11, functions that are transmitted across an abstraction barrier and then invoked across that same barrier in the opposite direction are called *callback functions*.

This simple model, however, is not sufficient to implement the other buttons you need for the Pathfinder application. The `quitAction` function, after all, is extremely simple and requires no information from the application. The functions that implement the other buttons, however, must have access to the graph representing the current map. The crux of the problem is therefore how to share information between the main program and the callback function triggered by an event.

One possible strategy—and indeed one that ends up being used in far too many user-interface packages—is to share information in global variables. This strategy, however, is a poor choice, because global variables can be manipulated in any part of the program. Such sharing is far too general. What we want instead is a mechanism that shares a data structure only between the cooperating parties: the main program that defines the button and the callback function invoked by `pathfinderEventLoop`. To accomplish this objective, the usual approach is to give the callback function an argument and pass this argument in both directions as a reference parameter.

The definition of `addButton` in `gpathfinder.h` implements this strategy by allowing an optional third argument that represents the data structure you want to share. The callback function must also then take an argument of that same type. When the user clicks on the button, the callback function receives this information and can therefore communicate with the main program without using global variables.

To illustrate this idea, suppose that your graph data structure is stored in a variable named `g` whose type is the extended `PathfinderGraph` class described in the description of Phase 1. You could then create a `Clear` button (not required for this assignment) by adding the line

```
addButton("Clear", clearAction, g);
```

to the main program. You could then define `clearAction` like this:

```
void clearAction(PathfinderGraph & g) {
    g.clear();
    drawPathfinderMap("");
    repaintPathfinderDisplay();
}
```

Note that the parameter `g` is passed by reference, so the call to `clear` actually clears the graph defined in the main program, and not some temporary copy of it.

The creation of the buttons on the screen requires just a few lines of code. The harder part is writing the callback functions. At this point, you don't yet have enough of the assignment in place to implement the `Dijkstra` and `Kruskal` buttons, but you can put them on the screen and assign them to functions whose bodies you will flesh out later. As part of Phase 4, however, you can (and indeed should) get the `Map` button working by having its callback function invoke the code you wrote in Phases 2 and 3.

Phase 5: Reimplement Dijkstra's algorithm so it fits the Pathfinder application

Figure 19-8 on page 823 of the reader has a complete implementation of Dijkstra's algorithm for finding the shortest path between two nodes. Your task in Phase 5 is simply to adapt that implementation so that it fits your data structure.

Copying the code for `findShortestPath` is not your primary task in Phase 5. The more challenging part is integrating the code into the application. When you click the **Dijkstra** button in the user interface, the program should ask the user to click on two nodes on the graph. These nodes then serve as the `start` and `finish` arguments to `findShortestPath`. More challenging still is the problem of highlighting the path after `findShortestPath` returns.

Phase 6: Design and implement the `Path` class and integrate it into the Dijkstra code

The code for Dijkstra's algorithm embodied in the `findShortestPath` function has some notable inefficiencies. One of these is that calculating the total distance along a path requires adding up the individual distances (or costs, in the language of the graph abstraction) of each of the arcs along the way. Because this operation must be performed several times for each path as it evolves, looping over the arcs in a path generates redundant computation that could easily be eliminated. Unfortunately, as long as the path is represented as a `vector` instead of as a separate class of its own, there is no obvious place to store the total distance alongside the list of arcs.

In Phase 6 of the assignment, your goal is to replace the `vector<Arc *>` used to store the path with a new `Path` class that you define in `path.h` and then implement in `path.cpp`. The versions of these files in the starter project are essentially empty. We don't tell you what methods the `Path` class must export or what instance variables must be in its private section. We do require that your class meet the following conditions:

1. The class must not export any public instance variables. All data required to store the path must be private to the class.
2. The method that returns the total cost of the path must run in constant time.
3. The class must export a `toString` method that returns a string composed of the nodes on the path separated by arrows formed by the two-character sequence `->`.
4. Any heap storage allocated by this class must be freed when the object is deleted.

Beyond those restrictions, you are free to design and implement this class as you see fit.

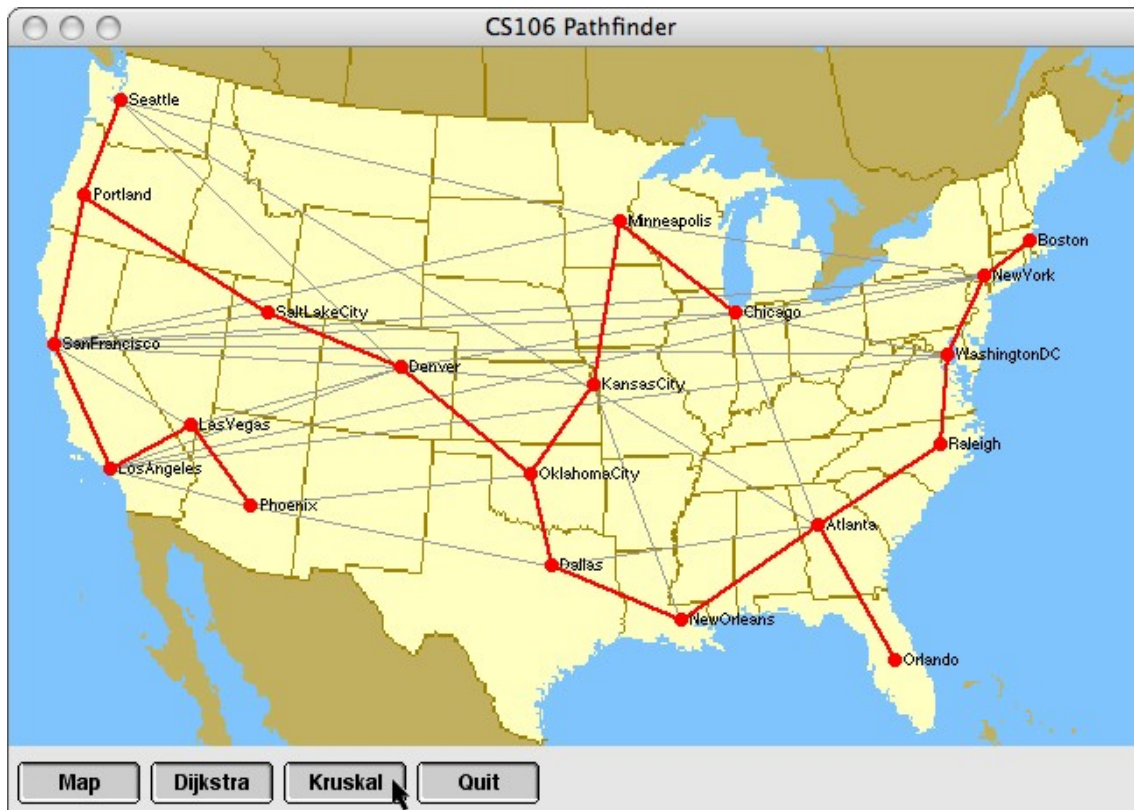
There is, however, one more requirement. Whenever you design a class, there are many different choices you could make in terms of what methods to export, what types to use, how flexible you want the interface to be, and so forth. Typically, each choice has advantages and disadvantages, so that the interface designer must evaluate the tradeoffs among the different possibilities. All too often, those design decisions are not reflected in the code, which means that future developers have to rediscover the underlying logic all over again. Thus, as the comments in the starter files make clear, one of the requirements of this assignment is that you include comments in the `path.cpp` file describing why you made the choices you did in your design of the `Path` class. This comment will be treated as an essay and will count for 25% of the style grade on the Pathfinder assignment.

Phase 7: Implement Kruskal's algorithm for finding minimum spanning trees

Your final task in the Pathfinder assignment is to implement another graph classic: Kruskal's algorithm for constructing a *minimum spanning tree* that connects the nodes in a graph at minimal cost. For example, if you invoke Kruskal's algorithm on the graph from the `USA.dat` file, you get the result shown in Figure 5 at the top of the next page.

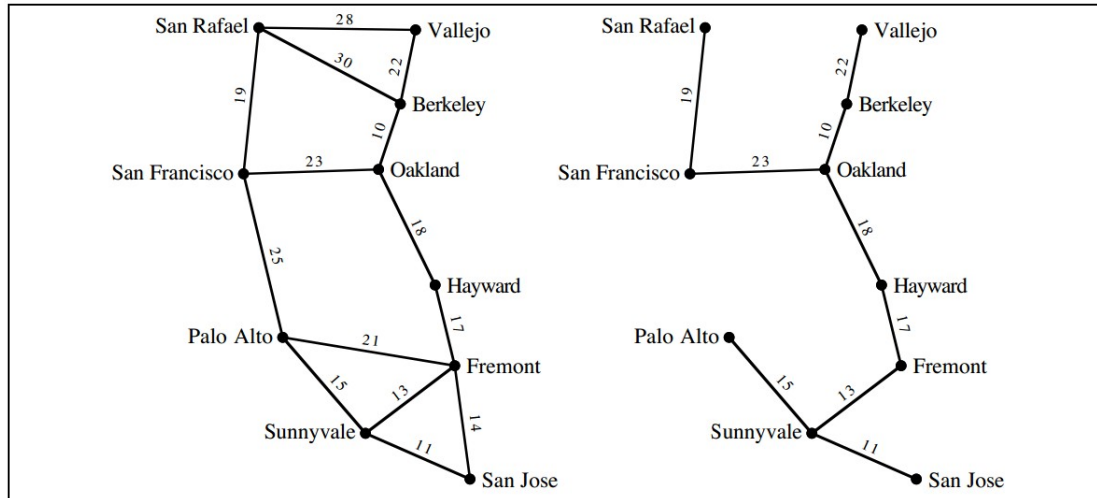
Although a description of Kruskal's algorithm appears in the reader starting on page 836, it seems worth repeating that discussion here so that you have all the information in one place.

Figure 5. Result of applying Kruskal's algorithm to generate a minimum spanning tree



In many situations, a minimum-cost path between two specific nodes is not as important as minimizing the cost of a network as a whole. As an example, suppose your company is building a new cable system that connects 10 large cities in the San Francisco Bay Area. Your preliminary research has provided you with cost estimates for laying new cable lines along a variety of possible routes. Those routes and their associated costs are shown in the graph on the left of Figure 6. Your job is to find the cheapest way to lay new cables so that all the cities are connected through some path.

Figure 6. A graph and its minimum spanning tree



To minimize the cost, one of the things you need to avoid is laying a cable that forms a cycle. Such a cable would be unnecessary, because some other path already links those cities, and thus you might as well leave such arcs out. The remaining graph forms a structure that is in some ways like a tree, even though it lacks a root node. Most importantly, the graph you're left with after eliminating the redundant arcs is tree-like in that it has no cycles. For historical reasons, this reduced graph that links all the nodes from the original one is called a *spanning tree*. The spanning tree in which the total cost associated with the arcs is as small as possible is called a *minimum spanning tree*. The cable-network problem is therefore equivalent to finding the minimum spanning tree of the graph, which is shown in the right side of Figure 6.

There are many algorithms in the literature for finding a minimum spanning tree. Of these, one of the simplest was devised by Joseph Kruskal in 1956. In Kruskal's algorithm, you consider the arcs in the graph in order of increasing cost. If the nodes at the endpoints of the arc are unconnected, then you include this arc as part of the spanning tree. If, however, the nodes are already connected by some path, you can discard this arc. The steps in the construction of the minimum spanning tree for the graph are shown in Figure 7, which was generated by a program that traces the operation of the algorithm.

Figure 7. Steps in the construction of the minimum spanning tree

```

Process edges in order of cost:
10: Berkeley -> Oakland
11: San Jose -> Sunnyvale
13: Fremont -> Sunnyvale
14: Fremont -> San Jose (not needed)
15: Palo Alto -> Sunnyvale
17: Fremont -> Hayward
18: Hayward -> Oakland
19: San Francisco -> San Rafael
21: Fremont -> Palo Alto (not needed)
22: Berkeley -> Vallejo
23: Oakland -> San Francisco
25: Palo Alto -> San Francisco (not needed)
28: San Rafael -> Vallejo (not needed)
30: Berkeley -> San Rafael (not needed)
    
```

Kruskal's is another example of a greedy algorithm. Since the goal is to minimize the total distance, it makes sense to consider shorter arcs before the longer ones. To process the arcs in order of increasing distance, the priority queue will come in handy again.

The tricky part of this algorithm is determining whether a given arc should be included. The strategy you will use is based on tracking connected sets. For each node, maintain the set of the nodes that are connected to it. At the start, each node is connected only to itself. When a new arc is added, you merge the sets of the two endpoints into one larger combined set that both nodes are now connected to. When considering an arc, if its two endpoints already belong to the same connected set, there is no point in adding that arc and thus you skip it. You continue considering arcs and merging connected sets until all nodes are joined into one set. The perfect data structure for tracking the connected sets is the `set` class, since it has the handy high-level operations (such as the union operation implemented as the `+` operator) that are exactly what you need here.

General hints and suggestions


- *Check out the demo.* Run our provided demo to learn how the program should operate.
- *Take care with your data structures.* Plan what data you need, where to store it, and how to access it. Think through the work to come and make sure your data structure adequately supports all your needs. Be sure you thoroughly understand the classes that you are using. Ask questions if anything is unclear.
- *Careful planning aids reuse.* This program has a lot of opportunity for unification and code reuse, but it requires some careful up-front planning. You'll find it much easier to do it right the first time than to go back and try to unify it after the fact. Sketch out your basic attack before writing any code and look for potential code-reuse opportunities in advance so you can design your functions to be all-purpose from the beginning.
- *Test on smaller data first.* There is a `Small.txt` data file with just four nodes that is helpful for early testing. The larger `USA.txt`, `Stanford.txt`, and `MiddleEarth.txt` data files are good for stress-testing once you have the basics in place.

Extension ideas

If you've completed a beautiful, elegant solution with time and energy to spare, there are many other graph explorations you could dive into, such as:

- *Implement the node and arc structure using classes instead of structures.* The hybrid design of the `Graph` class makes it possible to use `Node` and `Arc` classes instead of lower-level structures, as long as they export public instance variables with the appropriate names. In fact, if you simply replace the `Node` and `Arc` definitions in `graphtypes.h` with essentially identical versions of `Node` and `Arc` classes that export the same public fields, the program should continue to work as before. The advantage, however, of using classes here is that you can then define methods associated with nodes and arcs and then redesign the application to operate in a more object-oriented style.
- *Allow the user to find the shortest path according to different criteria.* If you wanted to fly from Portland to Atlanta, it would be tedious to take the minimum distance route shown in Figure 2, because that route requires you to take five different planes. The two-step route to Seattle and then directly to Atlanta is slightly longer in terms of distance, but likely to be shorter in terms of overall travel time. Extend the algorithm to allow the user to choose the number of segments (and possibly other criteria) to choose the optimal path.

- *Let the user trigger Dijkstra's algorithm by clicking on two nodes.* It's somewhat inconvenient to have to click the **Dijkstra** button before clicking on the nodes. Use the `defineClickListener` method in `gpathfinder.h` to implement a shortcut form in which the user simply clicks on the two nodes. The sample application includes this extension.
- *Employ A^* .* Dijkstra's algorithm is guaranteed to find the shortest path, but it can do unnecessary searching in what will prove to be the wrong direction if there are many short paths that lead away from the goal. One way to avoid this is to guide the search in the right direction by adding what computer scientists call a **heuristic**, which is a strategy that is likely to generate a better solution, although that outcome is not guaranteed. A heuristic is a "rule of thumb" that aids in the solution of a problem by employing domain-specific knowledge. You can change the shortest-path algorithm to include such domain-specific knowledge simply by changing how priorities are calculated. Instead of considering only the length of the path, add in an estimate of how close that path gets you to the goal. One good estimate is the straight-line distance from the end of the path to the goal. In Pathfinder, our understanding of the physical world tells us that Cartesian distance between two points is a good estimate of proximity (at least until distances get large enough that the distortions of actually being on the surface of a sphere become relevant). The search algorithm using this heuristic is an example of a group of algorithms called A^* in the artificial intelligence world. One of the inventors of A^* was Nils Nilsson, a Professor *emeritus* of our very own Stanford CS department. Add a button to run this heuristic and report on the difference in the numbers of paths considered between A^* and Dijkstra's algorithm.
- *Finding the graph center.* The city works department is thinking about adding a new fire station and need your help to find the appropriate central place. The goal is a position so that the shortest paths from the fire station to other locations are relatively small. The **eccentricity** of a node N is the maximum distance any other node is from N (*i.e.*, consider all of the shortest paths between N and each of the other nodes, the longest of those is N 's eccentricity). A node is at the **center** of the graph if its eccentricity is the smallest of all nodes (all nodes that tie for smallest are jointly considered the center). Add a feature to your program to find the graph center, which is the set of all nodes tied for the minimum eccentricity.
- *Max clique.* A graph **clique** is a subset of the graph nodes where every pair of nodes in the subset is directly connected. As in high school, it is a tight-knit cluster of friends that all know each other and ignore those outside their circle. Given the graph on the left below, the diagram on the right shows the maximal clique (the largest subset of the graph nodes that form a clique).



Identifying clusters of related objects often reduces to finding large cliques in graphs. For example, a program recently developed by the IRS to detect organized tax fraud looks for groups of phony tax

returns submitted in the hopes of getting undeserved refunds. The IRS constructs a graph where each tax form is a node and connects arcs between any two forms that appear suspiciously similar. A large clique in this graph points to fraud. Finding the truly maximal clique is known to require an exhaustive and time-consuming search, although there are good approximation algorithms that can find large cliques fairly efficiently. Add a feature to your program to find the maximal (or large) clique within the graph.

Deliverables

There is no interactive grading for this final assignment; all you have to do is submit your assignment in the normal way through paperless. The submission should include the four source files you've written or modified: `pathfinder.cpp`, `graphtypes.h`, `path.h`, and `path.cpp`. To save space on the submission area, you should not include the images we supplied.

Remember that we will not accept any late submissions for this assignment.