

Course Project: Graph Representation and Traversal

Functional Programming (ITI0212)

due: 2022-05-16

A *graph* is a data structure made up of *vertices* (or “nodes”) representing objects and *edges* representing relations between objects. There are several varieties of graph structures. For concreteness, in the following we will be discussing graphs in which,

- edges are directed, so an edge from x to y is distinct from an edge from y to x ,
- a given ordered pair of vertices may be connected by any number of edges,
- a vertex may be connected by edges to itself.

There are several possible ways to represent the data structure of a graph. Perhaps one of the simplest is as a list of edges, where each edge has a domain vertex, a codomain vertex, and a collection of attribute-value pairs. There is a standard language for representing graphs in this way called *DOT*. It is the graph description language for a program called *Graphviz*, which can render graph descriptions as diagrams. The specification of the DOT language can be found at <https://graphviz.org/doc/info/lang.html>.

For this project you will use Idris to implement a data structure for representing graphs, functions for reading and writing them from and to strings and files, an algorithm for efficiently traversing them, and something more of your choosing.

These tasks are intentionally less precisely specified than those in the labs and homeworks. You will need to use good judgement and programming practice to decompose each task into parts (types, functions, interfaces, etc.) that work together in a logical, modular, efficient, and elegant way to solve the problem at hand.

This leaves you with quite a lot of freedom, and you may take any approach you like provided that it conforms to the specifications and the spirit of the assignment. For example, it would violate the spirit of the assignment if you were to find an already implemented solution online, import it, and simply call the relevant functions. Thus, except for in the last task you must obtain permission from an instructor if you wish to use modules other than those found in the standard library.

Your solution to each of the following tasks should include a comment in your source file specifying the task number and briefly explaining the structure of your solution. If there are some parts of your program that are likely to be confusing, are unimplemented, or that contain known bugs, you should explain them in comments as well. In short, use good coding practice so that readers of your program (including your future self) can understand as easily as possible how it works. In any event, your submitted source file(s) should load without syntax or type errors.

Submit your project by pushing it to your course GitLab repository in a directory called **project**. Your submission should include a plain-text **README** file containing any information needed to load and run your program as intended. The marking criteria for the project is the same as for the homework assignments and is described in the course syllabus.

Task 1

Your first task is to write a function to read a graph description in a simplified version of the DOT language and construct the corresponding graph object if the description is well-formed:

```
read_graph : String -> Maybe Graph
```

You may represent a **Graph** any way you wish, but using a **List** of edges will probably be easiest at this stage. The simplified DOT language that you will need to support is for representing directed graphs using only edge specifiers. It is summarized here.

- The description should begin with the string literal **digraph** followed by an open brace { and end with a close brace }.
- Between the braces there should be a semicolon-separated sequence of edge specifiers.
- Each edge specifier should consist of an alphanumeric domain node name and codomain node name, separated by an ASCII arrow -> and followed by an optional attribute list.
- Each attribute list should be enclosed between brackets [and] and contain a comma-separated sequence of key-value pairs, themselves separated by an equals sign =.
- Whitespace, including tabs and linebreaks, is permitted but not required before or after any of the above-mentioned elements.

Here is a small example of a graph description in this simplified DOT language:

```
digraph
{
  a -> b [label = f , cost = 6] ;
  b -> c [label = g , cost = 7] ;
  c -> d [label = h , cost = 8]
}
```

You can see how Graphviz renders this graph by pasting it into the web site <https://edotor.net/>.

Your graph description reader does not need to be able to recognize and reject all invalid graph descriptions in the simplified DOT language. However, it should be able to recognize and accept valid graph descriptions.

Next, write a function that reads a simple graph description from a **.dot** file:

```
read_dot_file : (path : String) -> IO (Either FileError (Maybe Graph))
```

Task 2

Your next task is to write a function to take a graph and write a description of it in the simplified DOT language:

```
write_graph : Graph -> String
```

The graph descriptions that your writer emits should be valid Graphviz inputs. At this point you should test your reader and writer to ensure that if you read a valid graph description and then write the graph that you build from it, the result is itself a valid description of the same graph.

Next, write a function that writes a simple graph description to a **.dot** file:

```
write_dot_file : Graph -> (path : String) -> IO (Either FileError Unit)
```

Task 3

Your next task is to write a function that takes a graph and a designated “start” node and computes a subgraph of the input graph that includes just those edges that give a shortest path from the start node to each node in the input graph that is reachable from the start node.

```
shortest_paths : (start : Node) -> Graph -> Graph
```

This requires a bit of explanation.

A *subgraph* of a given graph consists of a subset of its nodes together with a subset of its edges, subject to the constraint that if an edge is included then so are its boundary nodes. When we represent graphs as edge lists this constraint is satisfied automatically.

A *path* in a graph is a finite sequence of consecutive edges, where two edges are *consecutive* if the codomain node of the first is the domain node of the second. For example, in the graph from task 1 we have a path from **a** to **d** given by **f ; g ; h**.

In order to have a notion of “shortest path” we need a notion of “edge length”. This is not necessarily intended to represent a geometric quantity, so we instead call it a *cost*, which may represent distance, time, currency, etc.. The cost of a path is the sum of the costs of its constituent edges. For example, the cost of the path **f ; g ; h** above is 21.

For what follows it is important that costs be non-negative. This is because we will need the cost of a path to be no greater than that of any path that extends it. For example, it should be the case that the cost of the path **f ; g ; h** is no greater than the cost of the path **f ; g ; h ; i**, regardless of what cost of the edge **i** is. In this task you may assume that the **cost** attribute of an edge specifier in a graph description is given as a natural number.

If an edge specifier does not include a **cost** attribute, or if the **cost** attribute of an edge cannot be parsed as a natural number, then we should consider the cost of that edge to be undefined. Only edges with well-defined costs should appear in the result of the **shortest_paths** function.

There is a simple algorithm for computing the **shortest_paths** function due to Dijkstra. Using the input graph, an “explored” subgraph of it, and a dictionary of node costs for the explored subgraph, we do the following.

- Determine an edge of the input graph whose domain node is explored (and thus has known cost), whose codomain node is unexplored (and thus has unknown cost), and such that the cost of the domain node plus the cost of the edge is minimal for edges from an explored node to an unexplored node.
- Add this edge to the explored subgraph and its codomain node to the cost dictionary, assigning it a cost of the sum of those of the edge and its domain node.

The process begins with the explored subgraph containing just the start node and the node cost dictionary assigning the start node a cost of 0. The process ends when we can no longer choose an edge with explored domain node, unexplored codomain node, and well-defined cost. At this point we just return the explored subgraph.

The algorithm works because if at some stage we choose an edge **e : x -> y**, it is because the cost of the node **x** plus the cost of the edge **e** is minimal for all edges from an explored node to an unexplored node. Thus the cost of getting to **y** using **e** can be no greater than the cost of getting to **y** any other way, and we have found a shortest path from the start node to **y**. This process terminates because the graph is finite and at each stage we reduce the size of the unexplored portion by one node.

Hint: in an imperative programming language we would probably implement this algorithm using a `while` loop. In Idris we don't have `while` loops, but we can achieve a similar effect using a higher-order (partial) function,

```
while_just : (f : a -> Maybe a) -> (x : a) -> a
```

which returns `x` if `f x` is `Nothing` (loop guard fails) and recurses on `y` if `f x` is `Just y` (loop guard succeeds).

Task 4

In this task you should choose some additional functionality to implement. Below are some possible options, which are meant to be inspirational, not normative. You may choose to do anything you like, provided that it is related to graphs and their representation or traversal.

This task will be evaluated the way that figure-skating is judged, in the sense that the ambition of the problem you choose and your success in solving it will both be taken into account. As a rule of thumb, to be eligible for full credit on this task you should do at least as much work as is needed for task 1. If you choose an ambitious problem and manage to solve it well, you can earn additional credit, making up for any marks lost in other parts of the project. However, in any case, you cannot earn more than a perfect score on the assignment as a whole.

If you choose to do something that augments the functionality from previous tasks, please make these changes as new functions so that your solution to task 4 can be evaluated independently from your solutions to tasks 1–3.

Highlighting the shortest paths: Write a function that takes a graph description and a designated start node and produces a graph description for the same graph but with the shortest paths from the start node highlighted in a different color. Make sure that Graphviz is able to render your output correctly.

Minimal spanning trees: Extend your solutions to tasks 1 and 2 to support undirected (a.k.a. bidirected) graphs, in which edges can be traversed in either direction. Then implement the *Prim algorithm* for finding a *minimal spanning tree* on a connected undirected graph.

Improved computational efficiency: The *Dijkstra algorithm* not only finds an efficient graph traversal but does so in a computationally efficient way. Re-implement your `shortest_paths` function in as computationally efficient a manner as you can. Use asymptotic analysis or testing to compare this implementation to the one you wrote in task 3. *Hint:* the `Data.SortedMap` module of the `contrib` library may be helpful.

A real parser: Parsing formal languages using the standard library functions from `Data.String` and `Data.List` is tedious and error-prone. In particular, it is quite difficult to detect all malformed expressions, such as those involving unbalanced parentheses. Higher-order functions provide a different, structural way to do parsing in the form of *parser combinators*. Re-implement your simple graph reader from task 1 using parser combinators so that it accepts *only* well-formed graph descriptions in the simplified DOT language. There are many good parser combinator tutorials to be found on the Internet. There is a simple implementation in the `Data.String.Parser` module of the `contrib` library.