

1 Trees

The trees (as discussed in computer science or mathematics) are either *undirected*, or *directed trees*, a.k.a. *rooted trees*. The notion of tree defined in Section 50 of the textbook is an *undirected tree*. Since many computer science applications use *directed trees* (or *rooted trees*) below we discuss this notion at length. Note that a directed tree is a special case of an undirected tree. This means that results proved for undirected trees hold also for directed trees but not vice versa.


1.1 Basic definitions

Recall that a graph is connected if there exists a walk between every pair of vertices within the graph. Further recall that a cycle is a walk that begins and ends in the same vertex and all other vertices of the walk are distinct. A graph is *acyclic* if it has no cycle.


Definition 1.1 (Undirected tree) *An undirected tree is a connected acyclic graph.*

In keeping with the arboreal nature of our terminology, a graph consisting of more than one connected component, where each connected component is a tree, is called a **forest**.

Example 1.1 *Each of the following graphs are trees.*



The following graph, containing three connected components, is a forest.



Often in applications trees are used to model hierarchical structures where the vertices have some kind of order (certain vertices are children of a given vertex etc.) Ordered trees are a restricted type of ordered graphs. We can define ordered trees by distinguishing a particular


vertex in an unordered tree as the *root* and giving direction to the edges so that they always “point away” from the root.

Definition 1.2 (Rooted tree) *A rooted tree is a tree with one vertex designated as the root. Every other edge in the tree is directed away from the root vertex.*

Note that since trees contain no cycles, any vertex has a unique path that connects it to the root (the vertex designated as the root). This means that the above notion “directed away from the root” is well defined.

When drawing rooted trees the convention is that the root is at the top. In drawing rooted trees we will sometimes omit the directions (i.e., the arrows) on each edge since both the root and the edge directions can be inferred from the way a tree is drawn.

Example 1.2 *The graph at left is a tree. The graph at right is the corresponding rooted tree obtained by choosing vertex a to be the root.*



1.2 Terminology for rooted trees


Since a tree is nothing more than a certain type of graph, all of the usual graph terminology applies to trees as well, such as “vertex” and “edge”. However, in addition to these terms, there exists a wealth of other terms relating specifically to trees.

We begin by summarizing some important terms relating to vertices and edges.

- The **root** of a tree is the unique vertex in a rooted tree with indegree zero.
- A **leaf** of a tree (also called an **external vertex**) is a vertex in a rooted tree with outdegree zero.
- An **internal vertex** is a vertex in a rooted tree with outdegree greater than zero.

- The **parent** of a vertex v is the unique vertex u incident to the directed edge (u, v) .
- A **child** of a vertex v is any vertex w that is adjacent to v by some directed edge (v, w) .
- Vertices that share a common parent vertex are called **siblings**.
- The **ancestors** of a vertex v are all vertices on the path from the root to v .
- The **descendants** of a vertex v are all vertices that have v as an ancestor.
- The **child degree** of a vertex v is equal to the number of children of v .

Example 1.3 Consider the tree T at left.




- The root of T is a .
- The leaves of T are d , e , f , and g .
- The internal vertices of T are a , b , and c .
- The parent of c is a . The children of c are e , f , and g . Each of these vertices are siblings.
- The ancestors of d are a and b .
- The only descendant of b is d .


We can also talk about the structure of the overall tree using specific terminology.

- Given a vertex v in a tree T , the **subtree** of T rooted at v is the subgraph containing v and its descendants.
- A tree T is an **m -ary tree** if each internal vertex of T has at most m children.
- An m -ary tree T is **full** if every internal vertex of T has exactly m children.

Example 1.4 Consider again the tree T from Example 1.3.



This tree is a 3-ary tree (also known as a **ternary tree**), since each internal vertex has at most three children. The tree is not full, however, because vertices *a* and *b* have fewer than three children. The following subtrees are contained within *T*:




The most common type of *m*-ary tree we will see is a **binary tree**, where each vertex of the tree has at most two children. Binary trees are prolific throughout computer science, so our notes will return to these trees from time to time.

Finally, there exist a few terms relating to the position of vertices within a tree.

- The **level** of a vertex *v* (also called the **depth** of *v*) is the length of the path from the root to *v*.
- The **height of a vertex** *v* is the length of the longest path from *v* to a leaf.
- The **height of a tree** *T* is the length of the longest path from the root of *T* to a leaf.

Example 1.5 Consider once more the tree *T* from Examples 1.3 and 1.4.



- Vertex *a* is on level 0 of *T*.
- Vertices *b* and *c* are on level 1 of *T*.
- Vertices *d*, *e*, *f*, and *g* are on level 2 of *T*.
- The height of vertex *a* is 2.
- The height of vertices *b* and *c* is 1.
- The height of *T* is 2.

1.3 Properties of Trees

Immediately from the definition of a tree, we can discover a number of properties about trees. For instance, all trees are bipartite graphs; we can illustrate this by colouring all vertices on odd levels of the tree with one colour and colouring all vertices on even levels of the tree

with another colour. Moreover, all trees are planar graphs.

Here, we prove some arguably more interesting properties of trees.

The following result gives a number of characterizations of undirected trees as special cases of a graph.

Theorem 1.1 *Given a graph T , the following statements are equivalent.*

1. T is an undirected tree.
2. There exists exactly one path between any pair of vertices in T .
3. T is connected and every edge of T is a cut edge.
4. T is acyclic and the addition of any edge to T creates a cycle.

Proof. To prove the equivalence of all of these statements, we will show that one statement implies the next statement in a chain of implications.

$1 \Rightarrow 2$: Let T be a tree, and suppose that there exists some pair of vertices u and v that are connected by two distinct paths, say p_1 and p_2 . We can reach v from u by following path p_1 , and we can return to u from v by following path p_2 . However, this implies that T contains a cycle, which contradicts our assumption that T is a tree. Therefore, no two vertices of T can be connected by more than one path.

$2 \Rightarrow 3$: Let T be a graph where every pair of vertices is connected by exactly one path. Then T is connected. Let e be an arbitrary edge of T , and let $T - e$ be the tree produced by removing e from T . If $T - e$ is connected, then there exists a path between the vertices that were previously incident to e . Since e itself was also a path between these vertices, then there must have existed more than one path between some pair of vertices in T , which contradicts our assumption. Therefore, $T - e$ is not connected and, since e is an arbitrary edge, any edge removed from T is a cut edge.

$3 \Rightarrow 4$: Let T be a connected graph where every edge is a cut edge. Suppose we add an edge e to T between two arbitrary vertices u and v . Since T is connected, a path already exists between u and v , so the addition of e to T creates a cycle between u and v . Since u and v

are arbitrary vertices, the addition of any edge to T creates a cycle.

4 \Rightarrow 1: Let T be an acyclic graph where the addition of any edge to T creates a cycle. To show that T is a tree, all we need to do is show that T is connected. Suppose that T is not connected. Then there exists some pair of vertices u and v such that no path exists between these vertices. However, this means that we can add an edge between u and v without creating a cycle, which contradicts our assumption. Therefore, T must be connected, and thus T is a tree. \square

There exist other equivalent characterizations of a tree beyond the four characterizations we proved here. For example, it is possible to show that a graph T is a tree if and only if T is connected *or* acyclic and T contains n vertices and $n - 1$ edges. Here, we will prove a simpler version of that statement on its own. First, however, we require a small intermediate result.

Lemma 1.1 *If T is a tree and vertex v of T has degree one, then $T - v$ is a tree.*

Proof. To show that $T - v$ is a tree, we must show that $T - v$ is both connected and acyclic.

We can see immediately that $T - v$ is acyclic, since if it were not, then our original tree T would contain a cycle. Since T is acyclic, $T - v$ must also be acyclic.

Next, consider two arbitrary vertices u and w from $T - v$. If $T - v$ is connected, then there exists a path between u and w . Moreover, this path cannot contain the removed vertex v ; if it did, then v would have to be located somewhere between u and w on the path, meaning that v would have degree 2. This contradicts our assumption. Therefore, for any pair of arbitrary vertices in $T - v$, there exists a path between those vertices that does not include v , and so $T - v$ is connected. \square

With this lemma, we can prove the aforementioned statement.

Theorem 1.2 *A tree T with n vertices contains $n - 1$ edges.*

Proof. We will prove the claim using the principle of mathematical induction on the number of vertices of T . Let $P(n)$ be the statement “a tree T with n vertices contains $n - 1$ edges”.

When $n = 1$, the tree contains one vertex and thus has no edges. Therefore, $P(1)$ is true.

Assume that $P(k)$ is true for some $k \in \mathbb{N}$.

We now show that $P(k + 1)$ is true; that is, a tree with $k + 1$ vertices contains k edges. Suppose v is a leaf of T . By Lemma 1.1, removing v from T produces a tree $T - v$ with k vertices. By our inductive hypothesis, $T - v$ contains $k - 1$ edges. It follows that T must contain k edges when we add v back into the tree.

Therefore, $P(k + 1)$ is true. By the principle of mathematical induction, $P(n)$ is true for all $n \in \mathbb{N}$. \square

1.4 Binary trees

Let us now shift from general trees to directed binary trees for a moment. Since binary trees are so popular in computer science, many results about these types of trees are known. For instance, given the height of a binary tree, we can determine certain measures such as the maximum number of leaves in that tree.

Theorem 1.3 *A binary tree T of height h contains at most 2^h leaves.*

Proof. We will prove the claim using the principle of mathematical induction on the height of T . Let $P(h)$ be the statement “a binary tree T of height h contains at most 2^h leaves”.

When $h = 1$, T consists of a root with at most 2 children, so T contains at most $2^1 = 2$ leaves. Therefore, $P(1)$ is true.

Assume that $P(h')$ is true for some $h' \in \mathbb{N}$.

We now show that $P(h' + 1)$ is true; that is, a binary tree of height $h' + 1$ contains at most $2^{h'+1}$ leaves. Let T be a binary tree of height $h' + 1$. Then T consists of a root with at most 2 children, both of which are the roots of subtrees of height at most h' . By our inductive hypothesis, both of these subtrees contain at most $2^{h'}$ leaves. Thus, T must contain at most $2 \times 2^{h'} = 2^{h'+1}$ leaves.

Therefore, $P(h' + 1)$ is true. By the principle of mathematical induction, $P(h)$ is true for all $h \in \mathbb{N}$. \square

Using a different proof technique, we can count the maximum number of vertices in a binary tree.

Theorem 1.4 *A binary tree T of height h contains at most $2^{h+1} - 1$ vertices.*

Proof. To prove this theorem, we count the maximum number of vertices at each level of the tree. Since we are dealing with a binary tree, the number of vertices at any given level may be at most twice the number of vertices at the previous level.

Suppose we have a binary tree T of height h . The zero'th level of T contains at most one vertex: the root. The first level contains at most twice the number of vertices at the zeroth level, so the first level contains at most two vertices. The second level contains at most four vertices, and so on.

Summing together each of these values, we get

$$\begin{aligned}\sum_{i=0}^h 2^i &= 2^0 + 2^1 + 2^2 + \cdots + 2^h \\ &= \frac{1 - 2^{h+1}}{1 - 2} \\ &= 2^{h+1} - 1.\end{aligned}$$

Thus, a binary tree of height h contains at most $2^{h+1} - 1$ vertices. \square


Both Theorems 1.3 and 1.4 can be generalized from binary trees to m -ary trees with the appropriate modifications. We will not present the generalizations here, but you should attempt to discover them yourself!

1.5 Representing Trees

Just like with graphs, there are a variety of different methods to represent trees. Since trees are a fundamental computer data structure, choosing the most appropriate way to represent a tree in memory or on a disk is vital if we want to optimize measures like access time or data retrieval.

The first class of tree representation methods use **linear** storage. Since rooted trees have a natural hierarchical property, we can represent the vertices of a tree quite easily in an array or a list. The array representation of a tree uses the parent/child relationship to establish at which index a vertex is to be located and allows for direct access to vertices. The list-of-lists method, on the other hand, relies heavily on recursion. The list-of-lists method is more similar to the adjacency list representation of a graph, but with additional nesting.

Example 1.6 Consider the following binary tree:




This tree can be stored in an array by storing first the root (labeled a), then the children of a (labeled b, c), then the children of b and the children of c . Alternatively, we may store the same tree in a list-of-lists representation as follows:


$$[a, [b, [d, [\emptyset], [\emptyset]], [e, [\emptyset], [\emptyset]]], [c, [\emptyset], [f, [\emptyset], [\emptyset]]]]$$

The second class of tree representation methods use **nonlinear** storage. Nonlinear storage not only gives us the advantage of not having to allocate a large block of contiguous memory, but it also has a similar structure to a tree itself: blocks of memory (vertices) connected by pointers (edges).

Example 1.7 Consider again the tree from Example 1.6:



This tree can be stored in a doubly-linked list (a form of nonlinear storage) as follows:



2 Tree Traversals

Now that we are familiar with trees and tree representations, we can consider methods of retrieving data stored in a tree. We call such methods **tree traversals**. Unlike arrays and lists, where we have one straightforward way to traverse the data within the structure, we can traverse data within a tree in more than one way. However, to maximize efficiency, we should strive to avoid visiting vertices more than once in a traversal; that is, we should only visit a vertex in a given tree exactly once during a traversal.

Tree traversal methods can be split into two categories: **depth-first** methods and **breadth-first** methods. Depth-first methods focus on traversing as deep within the tree as possible before moving to another part of the tree, while breadth-first methods visit all vertices on a given level before moving to a deeper level within the tree.

In what follows, we will only consider binary trees. However, each of the methods presented in this section can be generalized to m -ary trees.

We will also use a mnemonic system to remember the order of vertices visited for each tree traversal method. The letter V denotes the current vertex, the letter L denotes the left subtree of the current vertex, and the letter R denotes the right subtree of the current vertex.


2.1 Pre-order Traversal

In a pre-order traversal of a tree T , we visit the root of T first before descending into the left subtree and right subtree of the root, respectively. A pre-order traversal is a depth-first tree traversal, since we descend as far down into the left subtrees as possible before visiting the right subtrees. In our mnemonic system, the traversal order is VLR.

We can represent a pre-order traversal algorithm in pseudocode as follows:

```
preorder(T):
  if the root of T is empty:
    print nothing
  else:
    print the root of T
    preorder(left subtree of T)
    preorder(right subtree of T)
```

Example 2.1 Consider the following binary tree T :



A pre-order traversal of T gives the sequence of vertices $[a, b, d, g, h, e, c, f, i]$.

2.2 Post-order Traversal

A post-order traversal is, in a sense, the “opposite” of a pre-order traversal. In a post-order traversal of a tree T , we begin by visiting the left and right subtrees of the root of T before visiting the root itself. A post-order traversal is a depth-first traversal, by a similar line of reasoning as the pre-order traversal. In our mnemonic system, the traversal order is LRV.


We can represent a post-order traversal algorithm in pseudocode as follows:

```

postorder(T):
  if the root of T is empty:
    print nothing
  else:
    postorder(left subtree of T)
    postorder(right subtree of T)
    print the root of T

```

Example 2.2 Consider the same binary tree T from Example 2.1:



A post-order traversal of T gives the sequence of vertices $[g, h, d, e, b, i, f, c, a]$.

2.3 In-order Traversal

If we imagine the vertices of a binary tree as being ordered from left to right, where the leftmost vertex of the tree is the “first vertex” and the rightmost vertex of the tree is the “last vertex”, then we can develop a tree traversal method that preserves the left-to-right order of visiting vertices. The in-order traversal of a tree T goes as far left in the tree as possible before visiting a vertex. It then visits the parent of the left subtree before descending into the right subtree. Thus, an in-order traversal is a depth-first traversal. In our mnemonic system, the traversal order is LVR.

We can represent an in-order traversal algorithm in pseudocode as follows:

```

inorder(T):
  if the root of T is empty:
    print nothing
  else:


```

```

inorder(left subtree of T)
print the root of T
inorder(right subtree of T)

```

Example 2.3 Consider again the binary tree T from Examples 2.1 and 2.2:



An in-order traversal of T gives the sequence of vertices $[g, d, h, b, e, a, c, i, f]$.

2.4 Level-order Traversal

The level-order traversal is the only example of a breadth-first traversal we will see in this section. This method is a breadth-first traversal because it visits each vertex on a given level before descending further into the tree.

In order to keep track of vertices that we have yet to visit, we must use an auxiliary data structure. The natural choice in this scenario is a queue data structure, since we can add new vertices to the queue as we encounter them and remove vertices from the queue in the order in which they were added.


We can represent a level-order traversal algorithm in pseudocode as follows:

```

levelorder(T):
  add the root of T to queue Q
  while queue Q is not empty:
    dequeue vertex V from queue Q
    print V
    enqueue left child of V to queue Q
    enqueue right child of V to queue Q

```

Example 2.4 Consider once more the binary tree T from Examples 2.1, 2.2, and 2.3:



A level-order traversal of T gives the sequence of vertices $[a, b, c, d, e, f, g, h, i]$.

2.5 Applications of Traversals

Depending on what we wish to do with a tree, we may prefer to use one tree traversal method over another.

A pre-order traversal of a tree allows us to perform actions where we need to read the data in a parent vertex before the children vertices. For instance, given both a pre-order traversal and an in-order traversal, we may duplicate a tree. We can do so by comparing the positions of vertices in both traversal sequences.


Example 2.5 Suppose some binary tree has a pre-order traversal sequence $[a, b, d, e, c, f]$ and an in-order traversal sequence $[d, b, e, a, f, c]$. What is the structure of the binary tree?

First, observe that the initial element of the pre-order traversal sequence is the root of the tree. Thus, all other elements in the pre-order traversal sequence are children of the vertex a .

The in-order traversal sequence tells us in which subtree of the vertex a a given vertex belongs. Since vertices d , b , and e come before a in the in-order traversal sequence, each of these vertices are in the left subtree of a . Likewise, vertices f and c are in the right subtree of a .

From here, we can perform the same analysis on the “left subtree” subsequences $[b, d, e]$ and $[d, b, e]$, which reveal that b is the root of the left subtree, d is the left child of b , and e is the right child of b . Likewise, given the “right subtree” subsequences $[c, f]$ and $[f, c]$, we see that c is the root of the right subtree and f is the left child of c .

Altogether, we conclude that the structure of the binary tree is as follows:



A post-order traversal of a tree allows us to perform actions where we must operate on children vertices before the parent vertex. For instance, using a post-order traversal, we may delete a tree from the memory of a computer. It is vital that we delete vertices from the bottom-up, because if we delete a parent vertex before its children, then we lose references to the children vertices (and, therefore, leak memory that cannot be freed).

Finally, a **binary search tree** is a binary tree where each element is stored according to some order; for example, given some vertex u storing a value d , all values less than d are stored in vertices in the left subtree of u and all values greater than d are stored in vertices in the right subtree of u .

Unlike with general binary trees, it is possible to reconstruct a binary search tree using *only* the pre-order traversal sequence of the tree's vertices. We can do so by comparing the vertex to be inserted to the parent vertex. If the vertex contains a value smaller than that of the parent vertex, then we place it into the left subtree. If the vertex contains a value that is larger than that of the parent vertex, but smaller than that of the grandparent vertex, then we place the vertex to be inserted into the right subtree. Otherwise, we move up one level in the tree and check again.

We can also “collapse” the data stored in a binary search tree to one dimension by performing an in-order traversal. Doing so allows us to store the data (still in sorted order) in an array or list.

3 Spanning Trees

At this point, we know that trees are special kinds of graphs. But what can we say about the reverse relationship between these two structures; namely, how can we convert a graph into a tree? Of course, the answer seems clear: just delete enough edges from a graph to make it acyclic, but don't delete so many edges as to make the graph disconnected.


Naturally, this “clear” answer may leave you wanting. What steps must we take exactly to obtain a tree from a graph? And how do we know when we have done enough to obtain a tree? In this section, we investigate these questions and more, but first we must define the object of our study. The tree-from-a-graph that we will study here is known as a **spanning tree**.

Definition 3.1 (Spanning tree) *Given a connected graph G , a spanning tree of G is an unordered tree T where T is a subgraph of G and T includes every vertex in G .*

A spanning tree gets its name from the fact that it spans the vertices of a graph. Similarly, we may define **spanning forests** for a disconnected set of graphs, but we will not discuss that notion in as much depth here.

A spanning tree can be defined in terms of the vertices or the edges of a graph. In our definition, we focused on vertices. Alternatively, if we wish to focus on edges, then we say that a spanning tree is a maximal subset of edges of a graph that contains no cycle.

Example 3.1 Consider the following graph:



Some spanning trees contained in this graph are as follows. (This is not a complete list!)



Observe that we don't refer to a spanning tree of a graph as *the* spanning tree. This is because spanning trees need not be unique; a graph may contain many different spanning trees, and there are methods of counting the number of spanning trees a given graph contains. For general graphs, Kirchhoff's matrix-tree theorem, named for the German physicist Gustav Kirchhoff, allows us to calculate the number of spanning trees in a graph from the determinant of a matrix derived from the graph. (This theorem is a little outside the scope of this course, so we won't present it here.)

In addition to counting spanning trees themselves, we can count different measures on a spanning tree. For instance, as a consequence of Theorem 1.2, we know that any spanning tree in a graph with n vertices must contain exactly $n-1$ edges. Moreover, from Theorem 1.1, we know that adding any edge to a spanning tree will create a cycle. We call each of these cycles a **fundamental cycle**, and the study of such cycles leads to many more interesting graph theory results.

Before we continue, it is worth mentioning a small result that relates spanning trees to a certain property of their associated graph.

Theorem 3.1 A graph G is connected if and only if G contains a spanning tree.

Proof. (\Rightarrow): Suppose G is connected. If G is acyclic, then G is its own spanning tree and we are done. Otherwise, G contains at least one cycle. Select an arbitrary cycle in G and remove one edge e from the cycle. The resultant subgraph $G - e$ contains one fewer edge, but it remains connected. Repeat this process until no cycles remain. The resultant subgraph is a spanning tree of G .

(\Leftarrow): Suppose G contains a spanning tree T . By definition, T includes every vertex of G . Moreover, since T is a tree, we know by Theorem 1.1 that there exists exactly one path between every pair of vertices in T . Hence, there also exists a path between every pair of vertices in G , so G is connected. \square

3.1 Constructing Spanning Trees

Now comes the big question: how do we construct a spanning tree? We could use the technique given in the proof of Theorem 3.1, but the “taking away edges” approach is inefficient due to the fact that we must repeatedly identify cycles in the given graph. Instead, we can devise a “building up edges” approach. Just like with tree traversals, we have two different techniques we can use for this approach: we may construct a spanning tree of a graph in either a depth-first or a breadth-first manner.

In either case, we begin the construction process by selecting an arbitrary vertex v from our given graph G .

Using **depth-first search**, starting at vertex v , we build a path by visiting some vertex adjacent to v . We then continue this process of visiting new vertices and adding new edges to our path. If we reach some vertex and we cannot visit any other vertices (either because there are no more edges left to follow or because we have already visited all of the adjacent vertices), then we backtrack to the previous vertex and continue the process by branching off to a new path. We halt after we have visited every vertex of G .

We can use a stack data structure to keep track of the vertices we have visited as we traverse a path. To backtrack, we simply remove the top vertex from the stack.

We can represent a depth-first spanning tree construction method in pseudocode as follows:


DFSspanningtree(G):

```


T = a tree containing an arbitrary vertex V of G
push V to stack S
while S is not empty:
  for each vertex U adjacent to vertex at top of S:
    push U to stack S
    if U is already in T:
      pop U from stack S
    else:
      add U and incident edge to T
  if no more vertices are adjacent to vertex at top of S:
    pop vertex at top of S and backtrack

```

Example 3.2 Consider the following graph G :



Starting from vertex v in G , a depth-first search may construct a spanning tree in the following way:



Using **breadth-first search**, starting at vertex v , we add all vertices adjacent to v and edges incident to v to our tree T . We do not add a vertex if it is already in T . Then, we visit each adjacent vertex in the order in which it was added and repeat the process. Again, we halt after we have visited every vertex of G .

We can use a queue data structure to keep track of the vertices we have yet to visit as we

build our tree. The actions of enqueueing and dequeueing ensure that we visit vertices in the order in which they were added, as desired.

We can represent a breadth-first search spanning tree construction method in pseudocode as follows:

BFSspanningtree(G):

$T =$ a tree containing an arbitrary vertex V of G

enqueue V into queue Q

while queue Q is not empty:

 dequeue vertex from Q


 for each vertex U adjacent to current vertex:

 if U is not in T or Q :


 add U and incident edge to T

 enqueue U into queue Q

Example 3.3 Consider again the graph G from Example 3.2:



Starting from vertex v in G , a breadth-first search may construct a spanning tree in the following way:




3.2 Applications of Spanning Trees

Spanning trees find uses in a number of applications, both within and outside of computing. They are especially handy for applications where cycles within a graph structure are undesirable, since a spanning tree gives a subset of edges of a graph that retains every vertex while avoiding cycles. Thus, spanning trees can be used to create search engine crawlers, plan plumbing systems, lay out power grids, and optimize communication networks. For now, let's focus on the last application.

In a computer network, redundancy is useful, but too much redundancy can harm performance. We say that a computer network contains a “switching loop” when there exists more than one path between two endpoints in the network. Switching loops can provide resilience if, say, one path in the network becomes disconnected. However, since switches blindly send frames across all of their network connections, switching loops can impact data transmission if a frame gets caught in an infinite loop between switches. In the worst case, a large number of infinitely-looping frames can create “broadcast radiation” and grind normal network traffic to a halt.

As an example, the following computer network contains switching loops:



The Spanning Tree Protocol, invented by the American programmer Radia Perlman, uses spanning trees to avoid switching loops in a network. The protocol constructs a spanning tree across the network in much the same way as we constructed spanning trees across graphs, where network connections not included in the spanning tree are disabled. In this

way, switching loops are avoided and broadcast radiation is minimized. The protocol was later standardized by the IEEE and is still used to this day.


4 Minimum Spanning Trees

In this section we consider **weighted graphs** that are graphs in which each edge is assigned a numerical weight. These weights could represent distances, costs, capacities, or any other measure that might be applied to a graph. These weights are usually positive integer values, but they could also be zero, negative, or non-integer values.

Just like in unweighted graphs, we can find spanning trees in weighted graphs. Moreover, we can use exactly the same techniques to find spanning trees in weighted graphs as we did for unweighted graphs. However, since edge weights give us valuable information about the graph, we might wish to refine our methods of finding spanning trees to find bounds on the weights of edges in a spanning tree. If we focus on adding only the lowest-weighted edges to a spanning tree, then we get what is known as a **minimum spanning tree**.

Definition 4.1 (Minimum spanning tree) *Given a connected, weighted graph G , a minimum spanning tree of G is a spanning tree where the sum of weights of the edges in the tree is as small as possible.*

Example 4.1 *The weighted graph at left has a minimum spanning tree at right.*



It is possible to count the number of minimum spanning trees in a graph in a similar manner to how we counted general spanning trees. However, if we know one particular property of a graph in advance, then we get an interesting uniqueness result that does not apply to general spanning trees. This result tells us that if no pair of edges in the graph share the same weight, then the graph contains only one minimum spanning tree.

Theorem 4.1 *If a graph G is such that each edge in G has a distinct weight, then G contains a unique minimum spanning tree.*

Proof. Assume that a graph G with distinct-weight edges contains two minimum spanning trees T_1 and T_2 . Since these minimum spanning trees are different, there must exist at least one edge in one minimum spanning tree that does not exist in the other. Without loss of generality, let e_1 denote the edge of least weight that is in T_1 but not T_2 .

Since T_2 is a tree, adding e_1 to T_2 will create a cycle. Since both T_1 and T_2 were constructed from the same graph G , there must exist an edge e_2 in the cycle of T_2 that is not in T_1 . Moreover, e_2 must have a larger weight than e_1 . However, replacing e_2 with e_1 in T_2 would create a minimum spanning tree of smaller total weight, which contradicts our assumption that T_2 was a minimum spanning tree.

Therefore, there cannot exist more than one minimum spanning tree in G . \square

If we know other facts about certain aspects of a graph, then we can determine whether or not a given edge in the graph belongs to any minimum spanning tree of the graph. Our first lemma focuses on cut edges of a graph. (Recall that a cut edge is an edge that, when removed, increases the number of connected components of the graph.)

Lemma 4.1 (Cut property) *Given a graph G , if the weight of some cut edge e in G is smaller than the weight of any other cut edge in G , then e belongs to every minimum spanning tree of G .*

Proof. Assume that there exists a minimum spanning tree T of G that does not contain a cut edge e of smallest weight. Adding e to T will create a cycle that traverses the cut of G once via e and again via some other cut edge e' . If we remove the cut edge e' , then we would create a minimum spanning tree of smaller total weight, which contradicts our assumption that T was a minimum spanning tree.

Therefore, the cut edge e must belong to any minimum spanning tree T of G . \square

Our second lemma focuses on edges within a cycle of a graph. If we consider the weights of each edge in some cycle of a graph, then we can find at least one edge that cannot belong

to any minimum spanning tree of that graph.

Lemma 4.2 (Cycle property) *For any cycle C in a graph G , if the weight of some edge e in C is larger than the weight of any other edge in C , then e does not belong to any minimum spanning tree of G .*

Proof. Assume that there exists a minimum spanning tree T of G that contains such an edge e . By removing e from T , we disconnect the minimum spanning tree into two connected components. By adding any other unused edge in the cycle C , we can reconnect both connected components into one tree. Since all other edges in C are of smaller weight than e , adding any such edge would create a minimum spanning tree of smaller total weight, which contradicts our assumption that T was a minimum spanning tree.

Therefore, such an edge e cannot belong to any minimum spanning tree T of G . \square

4.1 Constructing Minimum Spanning Trees

Now, once again comes the big question: how do we construct a minimum spanning tree? Presumably, we don't want to use the same methods that we used for constructing general spanning trees, since those methods don't take edge weights into account. However, we can use the same basic ideas of selecting vertices and edges from a graph until we form a minimum spanning tree.

Our first method for constructing a minimum spanning tree, known as Kruskal's algorithm, works by sorting the edge set of a given graph G by increasing weight and then continually selecting the edge of least weight to add to the spanning tree. This method was invented by the American mathematician Joseph Kruskal in 1956.

Note that, although the final product is a connected minimum spanning tree, the intermediate trees need not be connected. Since Kruskal's algorithm selects edges only by weight, it may construct multiple connected components before linking them all together into one tree.


In pseudocode, Kruskal's algorithm is expressed as follows.

Kruskal(G):


```

T = an empty tree
sort all edges of G by increasing weight
for each edge  $e = \{u,v\}$ :
    if  $e$  does not create a cycle in T:
        add edge  $e$  to T
  
```

Example 4.2 Consider the following graph G :



Using Kruskal's algorithm, we get the following minimum spanning tree:



Although Kruskal's algorithm is intuitive and easy to understand, it comes with some downsides; for instance, before we even begin to build a spanning tree, we must first sort our edge set by increasing weight. For graphs with many edges, this process could take some time.

Our next method for constructing a minimum spanning tree takes a more vertex-centric approach, removing the requirement that we first sort the edge set. The Prim–Jarník algorithm works by selecting an arbitrary vertex v in a graph G , then constructing a spanning tree T by adding adjacent vertices by smallest edge weight. This algorithm uses the cut property of Lemma 4.1 to ensure that every edge added to T must belong to some minimum spanning tree of G .

The Prim–Jarník algorithm was originally invented by the Czech mathematician Vojtěch Jarník in 1930, and then rediscovered by the American mathematician Robert C. Prim in

1957—one year after Kruskal invented his method.

In pseudocode, the Prim–Jarník algorithm is expressed as follows.


PrimJarnik(G):

```


T = a tree containing an arbitrary vertex of G
while T contains fewer vertices than G:
    choose an edge  $e = \{u,v\}$  in  $G$  of smallest weight that is incident
    to a vertex  $v$  in  $T$ 
    if  $e$  does not create a cycle in  $T$ :
        add edge  $e$  and vertex  $v$  to  $T$ 

```

Example 4.3 Consider again the graph G from Example 4.2:



Using the Prim–Jarník algorithm starting from vertex v in G , we get the following minimum spanning tree:



Finally, as with so many other topics we have studied in this course, what we have seen here is not the complete picture. Other methods exist to construct minimum spanning trees. As an example, Borůvka's algorithm (invented by the Czech scientist Otakar Borůvka in 1926) was the first method to construct a minimum spanning tree. The algorithm works similarly

to the Prim–Jarník algorithm, but in parallel over the entire graph. Thus, it is especially useful if you are constructing minimum spanning trees on a computer that is capable of parallelization.