


## Graph Representation



Scientific results are mostly available in the form of articles in journals and conference proceedings, and on various Web<sup>1</sup> resources. These articles are not self-contained, but cite previous articles with related content. However, when you read an article from 1975 with an interesting partial result, you may often ask yourself what the current state of the art is. In particular, you would like to know which newer articles cite the old article. Projects such as Google Scholar provide this functionality by analyzing the reference sections of articles and building a database of articles that efficiently supports looking up articles that cite a given article.

We can easily model this situation by a directed graph. The graph has a node for each article and an edge for each citation. An edge  $(u, v)$  from article  $u$  to article  $v$  means that  $u$  cites  $v$ . In this terminology, every node (= article) stores all its outgoing edges (= the articles cited by it) but not the incoming edges (the articles citing it). If every node were also to store the incoming edges, it would be easy to find the citing articles. One of the main tasks of Google Scholar is to construct the reversed edges. This example shows that the cost of even a very basic elementary operation on a graph, namely finding all edges entering a particular node, depends heavily on the representation of the graph. If the incoming edges are stored explicitly, the operation is easy; if the incoming edges are not stored, the operation is nontrivial.

In this chapter, we shall give an introduction to the various possibilities for representing graphs in a computer. We focus mostly on directed graphs and assume that an undirected graph  $G = (V, E)$  is represented as the corresponding (bi)directed graph  $G' = (V, \cup_{\{u,v\} \in E} \{(u,v), (v,u)\})$ . Figure 8.1 illustrates the concept of a bidirected graph. Most of the data structures presented also allow us to represent multiple parallel edges and self-loops. We start with a survey of the operations that we may want to support.

- *Accessing associated information.* Given a node or an edge, we frequently want to access information associated with it, for example the weight of an edge or the distance to a node. In many representations, nodes and edges are objects, and we can store this information directly as a member of these objects. If not otherwise mentioned, we assume that  $V = 1..n$  so that information associated

<sup>1</sup> The picture above shows a spider web (USFWS, see [http://commons.wikimedia.org/wiki/Image:Water\\_drops\\_on\\_spider\\_web.jpg](http://commons.wikimedia.org/wiki/Image:Water_drops_on_spider_web.jpg)).

with nodes can be stored in arrays. When all else fails, we can always store node or edge information in a hash table. Hence, accesses can be implemented to run in constant time. In the remainder of this book we abstract from the various options for realizing access by using the data types *NodeArray* and *EdgeArray* to indicate array-like data structures that can be indexed by nodes and by edges, respectively.

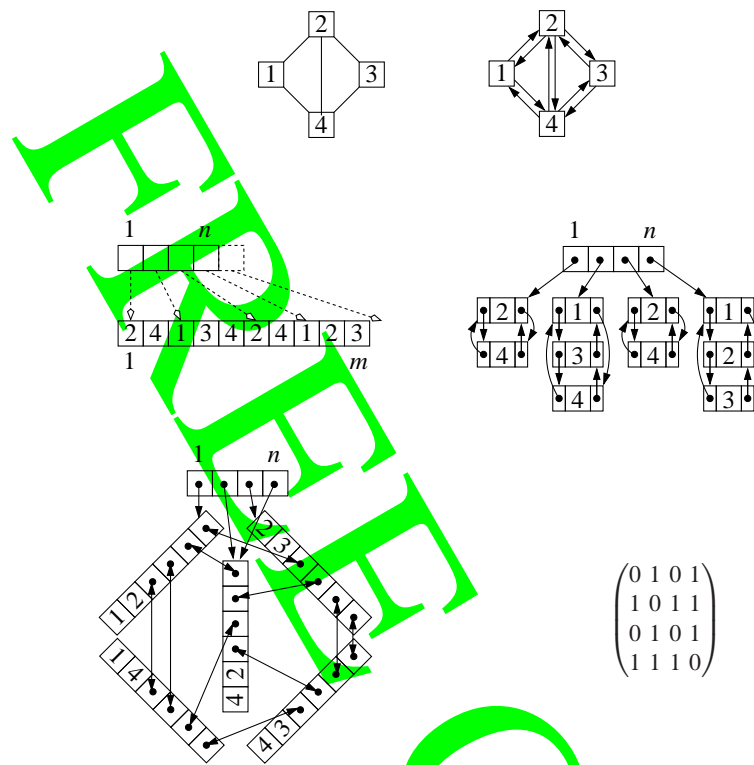
- *Navigation.* Given a node, we may want to access its outgoing edges. It turns out that this operation is at the heart of most graph algorithms. As we have seen in the example above, we sometimes also want to know the incoming edges.
- *Edge queries.* Given a pair of nodes  $(u, v)$ , we may want to know whether this edge is in the graph. This can always be implemented using a hash table, but we may want to have something even faster. A more specialized but important query is to find the *reverse edge*  $(v, u)$  of a directed edge  $(u, v) \in E$  if it exists. This operation can be implemented by storing additional pointers connecting edges with their reversals.
- *Construction, conversion and output.* The representation most suitable for the algorithmic problem at hand is not always the representation given initially. This is not a big problem, since most graph representations can be translated into each other in linear time.
- *Update.* Sometimes we want to add or remove nodes or edges. For example, the description of some algorithms is simplified if a node is added from which all other nodes can be reached (e.g. Fig. 10.10).

## 8.1 Unordered Edge Sequences

Perhaps the simplest representation of a graph is as an unordered sequence of edges. Each edge contains a pair of node indices and, possibly, associated information such as an edge weight. Whether these node pairs represent directed or undirected edges is merely a matter of interpretation. Sequence representation is often used for input and output. It is easy to add edges or nodes in constant time. However, many other operations, in particular navigation, take time  $\Theta(m)$ , which is forbiddingly slow. Only a few graph algorithms work well with the edge sequence representation; most algorithms require easy access to the edges incident on any given node. In this case the ordered representations discussed in the following sections are appropriate. In Chap. 11, we shall see two minimum-spanning-tree algorithms: one works well with an edge sequence representation and the other needs a more sophisticated data structure.

## 8.2 Adjacency Arrays – Static Graphs

To support easy access to the edges leaving any particular node, we can store the edges leaving any node in an array. If no additional information is stored with the edges, this array will just contain the indices of the target nodes. If the graph is *static*, i.e., does not change over time, we can concatenate all these little arrays into a single



**Fig. 8.1.** The first row shows an undirected graph and the corresponding bidirected graph. The second row shows the adjacency array and adjacency list representations of this bidirected graph. The third row shows the linked-edge-objects representation and the adjacency matrix

edge array  $E$ . An additional array  $V$  stores the starting positions of the subarrays, i.e., for any node  $v$ ,  $V[v]$  is the index in  $E$  of the first edge out of  $v$ . It is convenient to add a dummy entry  $V[n + 1]$  with  $V[n + 1] = m + 1$ . The edges out of any node  $v$  are then easily accessible as  $E[V[v]]$ ,  $\dots$ ,  $E[V[v + 1] - 1]$ ; the dummy entry ensures that this also holds true for node  $n$ . Figure 8.1 shows an example.

The memory consumption for storing a directed graph using adjacency arrays is  $n + m + \Theta(1)$  words. This is even more compact than the  $2m$  words needed for an edge sequence representation.

Adjacency array representations can be generalized to store additional information: we may store information associated with edges in separate arrays or within the edge array. If we also need incoming edges, we may use additional arrays  $V'$  and  $E'$  to store the reversed graph.

**Exercise 8.1.** Design a linear-time algorithm for converting an edge sequence representation of a directed graph into an adjacency array representation. You should use

only  $O(1)$  auxiliary space. Hint: view the problem as the task of sorting edges by their source node and adapt the integer-sorting algorithm shown in Fig. 5.15.

### 8.3 Adjacency Lists – Dynamic Graphs

Edge arrays are a compact and efficient graph representation. Their main disadvantage is that it is expensive to add or remove edges. For example, assume that we want to insert a new edge  $(u, v)$ . Even if there is room in the edge array  $E$  to accommodate it, we still have to move the edges associated with nodes  $u + 1$  to  $n$  one position to the right, which takes time  $O(m)$ .

In Chap. 3, we learned how to implement dynamic sequences. We can use any of the solutions presented there to produce a dynamic graph data structure. For each node  $v$ , we represent the sequence  $E_v$  of outgoing (or incoming, or outgoing and incoming) edges by an unbounded array or by a (singly or doubly) linked list. We inherit the advantages and disadvantages of the respective sequence representations. Unbounded arrays are more cache-efficient. Linked lists allow constant-time insertion and deletion of edges at arbitrary positions. Most graphs arising in practice are sparse in the sense that every node has only a few incident edges. Adjacency lists for sparse graphs should be implemented without the dummy item introduced in Sect. 3.1, because an additional item would waste considerable space. In the example in Fig. 8.1, we show circularly linked lists.

**Exercise 8.2.** Suppose the edges adjacent to a node  $u$  are stored in an unbounded array  $E_u$ , and an edge  $e = (u, v)$  is specified by giving its position in  $E_u$ . Explain how to remove  $e = (u, v)$  in constant amortized time. Hint: you do *not* have to maintain the relative order of the other edges.

**Exercise 8.3.** Explain how to implement the algorithm for testing whether a graph is acyclic discussed in Chap. 2.9 so that it runs in linear time, i.e., design an appropriate graph representation and an algorithm using it efficiently. Hint: maintain a queue of nodes with outdegree zero.

Bidirected graphs arise frequently. Undirected graphs are naturally presented as bidirected graphs, and some algorithms that operate on directed graphs need access not only to outgoing edges but also to incoming edges. In these situations, we frequently want to store the information associated with an undirected edge or a directed edge and its reversal only once. Also, we may want to have easy access from an edge to its reversal.

We shall describe two solutions. The first solution simply associates two additional pointers with every directed edge. One points to the reversal, and the other points to the information associated with the edge.

The second solution has only one item for each undirected edge (or pair of directed edges) and makes this item a member of two adjacency lists. So, the item for an undirected edge  $\{u, v\}$  would be a member of lists  $E_u$  and  $E_v$ . If we want doubly

linked adjacency information, the edge object for any edge  $\{u, v\}$  stores four pointers: two are used for the doubly linked list representing  $E_u$ , and two are used for the doubly linked list representing  $E_v$ . Any node stores a pointer to some edge incident on it. Starting from it, all edges incident on the node can be traversed. The bottom part of Fig. 8.1 gives an example. A small complication lies in the fact that finding the other end of an edge now requires some work. Note that the edge object for an edge  $\{u, v\}$  stores the endpoints in no particular order. Hence, when we explore the edges out of a node  $u$ , we must inspect both endpoints and then choose the one which is different from  $u$ . An elegant alternative is to store  $u \oplus v$  in the edge object [145]. An exclusive OR with either endpoint then yields the other endpoint. Also, this representation saves space.

## 8.4 The Adjacency Matrix Representation

An  $n$ -node graph can be represented by an  $n \times n$  adjacency matrix  $A$ .  $A_{ij}$  is 1 if  $(i, j) \in E$  and 0 otherwise. Edge insertion or removal and edge queries work in constant time. It takes time  $O(n)$  to obtain the edges entering or leaving a node. This is only efficient for very dense graphs with  $m = \Omega(n^2)$ . The storage requirement is  $n^2$  bits. For very dense graphs, this may be better than the  $n + m + O(1)$  words required for adjacency arrays. However, even for dense graphs, the advantage is small if additional edge information is needed.

**Exercise 8.4.** Explain how to represent an undirected graph with  $n$  nodes and without self-loops using  $n(n-1)/2$  bits.

Perhaps more important than actually storing the adjacency matrix is the conceptual link between graphs and linear algebra introduced by the adjacency matrix. On the one hand, graph-theoretic problems can be solved using methods from linear algebra. For example, if  $C = A^k$ , then  $C_{ij}$  counts the number of paths from  $i$  to  $j$  with exactly  $k$  edges.

**Exercise 8.5.** Explain how to store an  $n \times n$  matrix  $A$  with  $m$  nonzero entries using storage  $O(m+n)$  such that a matrix-vector multiplication  $Ax$  can be performed in time  $O(m+n)$ . Describe the multiplication algorithm. Expand your representation so that products of the form  $x^T A$  can also be computed in time  $O(m+n)$ .

On the other hand, graph-theoretic concepts can be useful for solving problems from linear algebra. For example, suppose we want to solve the matrix equation  $Bx = c$ , where  $B$  is a symmetric matrix. Now consider the corresponding adjacency matrix  $A$  where  $A_{ij} = 1$  if and only if  $B_{ij} \neq 0$ . If an algorithm for computing connected components finds that the undirected graph represented by  $A$  contains two distinct connected components, this information can be used to reorder the rows and columns of  $B$  such that we obtain an equivalent equation of the form

$$\begin{pmatrix} B_1 & 0 \\ 0 & B_2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}.$$

This equation can now be solved by solving  $B_1x_1 = c_1$  and  $B_2x_2 = c_2$  separately. In practice, the situation is more complicated, since we rarely have matrices whose corresponding graphs are disconnected. Still, more sophisticated graph-theoretic concepts such as cuts can help to discover structure in the matrix which can then be exploited in solving problems in linear algebra.

## 8.5 Implicit Representations

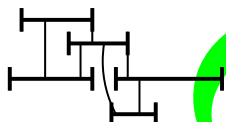
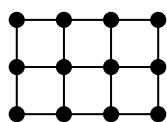
Many applications work with graphs of special structure. Frequently, this structure can be exploited to obtain simpler and more efficient representations. We shall give two examples.

The *grid graph*  $G_{k\ell}$  with node set  $V = [0..k-1] \times [0..\ell-1]$  and edge set

$$E = \{((i, j), (i, j')) \in V^2 : |j - j'| = 1\} \cup \{((i, j), (i', j)) \in V^2 : |i - i'| = 1\}$$

is completely defined by the two parameters  $k$  and  $\ell$ . Figure 8.2 shows  $G_{3,4}$ . Edge weights could be stored in two two-dimensional arrays, one for the vertical edges and one for the horizontal edges.

An *interval graph* is defined by a set of intervals. For each interval, we have a node in the graph, and two nodes are adjacent if the corresponding intervals overlap.



**Fig. 8.2.** The grid graph  $G_{3,4}$  (left) and an interval graph with five nodes and six edges (right)

### Exercise 8.6 (representation of interval graphs).

- Show that for any set of  $n$  intervals there is a set of intervals whose endpoints are integers in  $[1..2n]$  and that defines the same graph.
- Devise an algorithm that decides whether the graph defined by a set of  $n$  intervals is connected. Hint: sort the endpoints of the intervals and then scan over the endpoints in sorted order. Keep track of the number of intervals that have started but not ended.
- (c\*) Devise a representation for interval graphs that needs  $O(n)$  space and supports efficient navigation. Given an interval  $I$ , you need to find all intervals  $I'$  intersecting it;  $I'$  intersects  $I$  if  $I$  contains an endpoint of  $I'$  or  $I \subseteq I'$ . How can you find the former and the latter kinds of interval?

## 8.6 Implementation Notes

We have seen several representations of graphs in this chapter. They are suitable for different sets of operations on graphs, and can be tuned further for maximum

performance in any particular application. The edge sequence representation is good only in specialized situations. Adjacency matrices are good for rather dense graphs. Adjacency lists are good if the graph changes frequently. Very often, some variant of adjacency arrays is fastest. This may be true even if the graph changes, because often there are only a few changes, or all changes happen in an initialization phase of a graph algorithm, or changes can be agglomerated into occasional rebuildings of the graph, or changes can be simulated by building several related graphs.

There are many variants of the adjacency array representation. Information associated with nodes and edges may be stored together with these objects or in separate arrays. A rule of thumb is that information that is frequently accessed should be stored with the nodes and edges. Rarely used data should be kept in separate arrays, because otherwise it would often be moved to the cache without being used. However, there can be other, more complicated reasons why separate arrays may be faster. For example, if both adjacency information and edge weights are read but only the weights are changed, then separate arrays may be faster because the amount of data written back to the main memory is reduced.

Unfortunately, no graph representation is best for all purposes. How can one cope with the zoo of graph representations? First, libraries such as LEDA and the Boost graph library offer several different graph data types, and one of them may suit your purposes. Second, if your application is not particularly time- or space-critical, several representations might do and there is no need to devise a custom-built representation for the particular application. Third, we recommend that graph algorithms should be written in the style of generic programming [71]. The algorithms should access the graph data structure only through a small set of operations, such as iterating over the edges out of a node, accessing information associated with an edge, and proceeding to the target node of an edge. The interface can be captured in an interface description, and a graph algorithm can be run on any representation that realizes the interface. In this way, one can experiment with different representations. Fourth, if you have to build a custom representation for your application, make it available to others.

### 8.6.1 C++

LEDA [131, 118, 145] offers a powerful graph data type that supports a large variety of operations in constant time and is convenient to use, but is also space-consuming. Therefore LEDA also implements several more space-efficient adjacency array representations.

The Boost graph library [27, 119] emphasizes a strict separation of representation and interface. In particular, Boost graph algorithms run on any representation that realizes the Boost interface. Boost also offers its own graph representation class *adjacency\_list*. A large number of parameters allow one to choose between variants of graphs (directed and undirected graphs and multigraphs<sup>2</sup>), types of navigation available (in-edges, out-edges, . . .), and representations of vertex and edge sequences

<sup>2</sup> Multigraphs allow multiple parallel edges.



(arrays, linked lists, sorted sequences, . . .). However, it should be noted that the array representation uses a separate array for the edges adjacent to each vertex.

### 8.6.2 Java

JDSL [78] offers rich support for graphs in *jdsl.graph*. It has a clear separation between interfaces, algorithms, and representation. It offers an adjacency list representation of graphs that supports directed and undirected edges.

## 8.7 Historical Notes and Further Findings

Special classes of graphs may result in additional requirements for their representation. An important example is *planar graphs* – graphs that can be drawn in the plane without edges crossing. Here, the ordering of the edges adjacent to a node should be in counterclockwise order with respect to a planar drawing of the graph. In addition, the graph data structure should efficiently support iterating over the edges along a *face* of the graph, a cycle that does not enclose any other node. LEDA offers representations for planar graphs.

Recall that *bipartite graphs* are special graphs where the node set  $V = L \cup R$  can be decomposed into two disjoint subsets  $L$  and  $R$  such that the edges are only between nodes in  $L$  and  $R$ . All representations discussed here also apply to bipartite graphs. In addition, one may want to store the two sides  $L$  and  $R$  of the graph.

*Hypergraphs*  $H = (V, E)$  are generalizations of graphs, where edges can connect more than two nodes. Hypergraphs are conveniently represented as the corresponding bipartite graph  $B_H = (E \cup V, \{(e, v) : e \in E, v \in V, v \in e\})$ .

*Cayley graphs* are an interesting example of implicitly defined graphs. Recall that a set  $V$  is a *group* if it has an associative multiplication operation  $*$ , a neutral element, and a multiplicative inverse operation. The *Cayley graph*  $(V, E)$  with respect to a set  $S \subseteq V$  has the edge set  $\{(u, u * s) : u \in V, s \in S\}$ . Cayley graphs are useful because graph-theoretic concepts can be useful in group theory. On the other hand, group theory yields concise definitions of many graphs with interesting properties. For example, Cayley graphs have been proposed as interconnection networks for parallel computers [12].

In this book, we have concentrated on convenient data structures for *processing* graphs. There is also a lot of work on *storing* graphs in a flexible, portable, space-efficient way. Significant compression is possible if we have a priori information about the graphs. For example, the edges of a triangulation of  $n$  points in the plane can be represented with about  $6n$  bits [42, 168].