Lecture 1

First Steps in Graph Theory

This lecture introduces Graph Theory, the main subject of the course, and includes some basic definitions as well as a number of standard examples.

Reading: Some of the material in today's lecture comes from the beginning of Chapter 1 in

> Dieter Jungnickel (2013), *Graphs, Networks and Algorithms*, 4th edition, which is available online via SpringerLink.

If you are at the university, either physically or via the VPN, you can download the chapters of this book as PDFs.

1.1 The Königsberg Bridge Problem

Graph theory is usually said to have been invented in 1736 by the great Leonhard Euler, who used it to solve the Königsberg Bridge Problem. I used to find this hard to believe—the graph-theoretic graph is such a natural and useful abstraction that it's difficult to imagine that no one hit on it earlier—but Euler's paper about graphs[1] is generally acknowledged[2] as the first one and it certainly provides a satisfying solution to the bridge problem. The sketch in the left panel of Figure 1.1 comes from Euler's original paper and shows the main features of the problem. As one can see by comparing Figures 1.1 and 1.2, even this sketch is already a bit of an abstraction.

The question is, can one make a walking tour of the city that (a) starts and finishes in the same place and (b) crosses every bridge exactly once. The short answer to this question is "No" and the key idea behind proving this is illustrated in the right panel of Figure 1.1. It doesn't matter what route one takes while walking around on, say, the smaller island: all that really matters are the ways in which the bridges connect the four land masses. Thus we can shrink the small island to a

[1] L. Euler (1736), Solutio problematis ad geometriam situs pertinentis, *Commentarii Academiae Scientiarum Imperialis Petropolitanae* **8**, pp. 128–140.

[2] See, for example, Robin Wilson and John J. Watkins (2013), *Combinatorics: Ancient & Modern*, OUP. ISBN 978-0-19-965659-2.

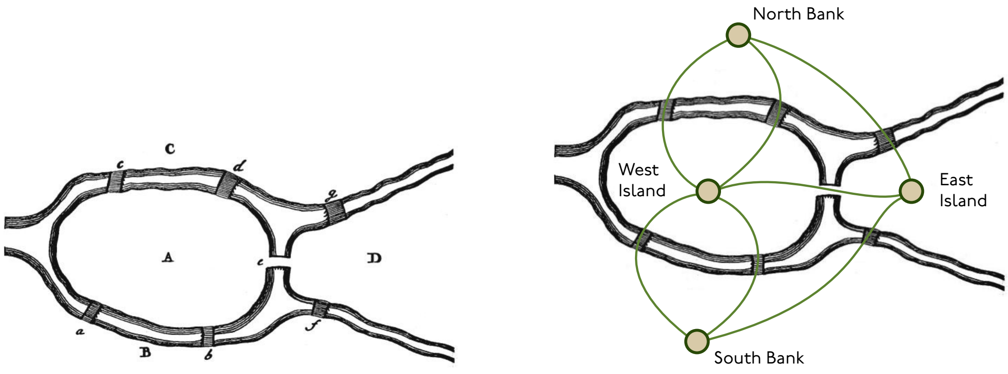

Figure 1.1: *The panel at left shows the seven bridges and four land masses that provide the setting for the Königsberg bridge problem, which asks whether it is possible to make a circular walking tour of the city that crosses every bridge exactly once. The panel at right includes a graph-theoretic abstraction that helps one prove that no such tour exists.*

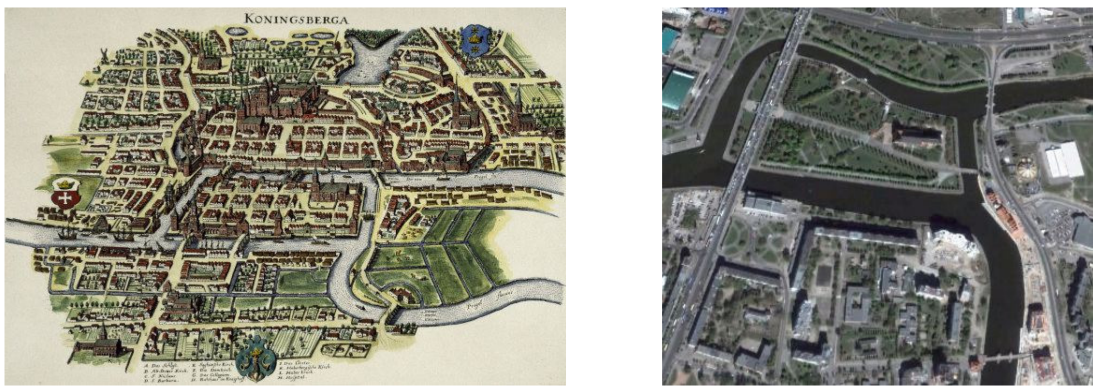

Figure 1.2: *Königsberg is a real place—a port on the Baltic—and during Euler's lifetime it was part of the Kingdom of Prussia. The panel at left is a bird's-eye view of the city that shows the celebrated seven bridges. It was made by Matthäus Merian and published in 1652. The city is now called Kaliningrad and is part of the Russian Federation. It was bombed heavily during the Second World War: the panel at right shows a recent satellite photograph and one can still recognize the two islands and modern versions of some of the bridges, but very little else appears to remain.*

point—and do the same with the other island, as well as with the north and south banks of the river—and then connect them with arcs that represent the bridges. The problem then reduces to the question whether it is possible to draw a path that starts and finishes at the same dot, but traces each of over the seven arcs exactly once.

One can prove that such a tour is impossible by contradiction. Suppose that one exists: it must then visit the easternmost island (see Figure 1.3) and we are free to imagine that the tour actually starts there. To continue we must leave the island, crossing one of its three bridges. Then, later, because we are required to

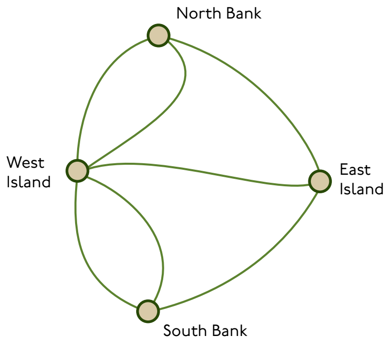

Figure 1.3: *The Königsberg Bridge graph on its own: it is not possible to trace a path that starts and ends on the eastern island without crossing some bridge at least twice.*

cross each bridge exactly once, we will have to return to the eastern island via a different bridge from the one we used when setting out. Finally, having returned to the eastern island once, we will need to leave again in order to cross the island's third bridge. But then we will be unable to return without recrossing one of the three bridges. And this provides a contradiction: the walk is supposed to start and finish in the same place and cross each bridge exactly once.

1.2 Definitions: graphs, vertices and edges

The abstraction behind Figure 1.3 turns out to be very powerful: one can draw similar diagrams to represent "connections" between "things" in a very general way. Examples include: representations of social networks in which the points are people and the arcs represent acquaintance; genetic regulatory networks in which the points are genes and the arcs represent activation or repression of one gene by another and scheduling problems in which the points are tasks that contribute to some large project and the arcs represent interdependence among the tasks. To help us make more rigorous statements, we'll use the following definition:

Definition 1.1. *A **graph** is a finite, nonempty set V, the **vertex set**, along with a set E, the **edge set**, whose elements $e \in E$ are pairs $e = (a, b)$ with $a, b \in V$.*

We will often write $G(V, E)$ to mean the graph G with vertex set V and edge set E. An element $v \in V$ is called a *vertex* (plural *vertices*) while an element $e \in E$ is called an *edge*.

The definition above is deliberately vague about whether the pairs that make up the edge set E are ordered pairs—in which case (a, b) and (b, a) with $a \neq b$ are distinct edges—or unordered pairs. In the unordered case (a, b) and (b, a) are just two equivalent ways of representing the same pair.

Definition 1.2. *An **undirected graph** is a graph in which the edge set consists of unordered pairs.*

Figure 1.4: *Diagrams representing graphs with vertex set $V = \{a, b\}$ and edge set $E = \{(a, b)\}$. The diagram at left is for an undirected graph, while the one at right shows a directed graph. Thus the arrow on the right represents the ordered pair (a, b).*

Definition 1.3. *A* ***directed graph*** *is a graph in which the edge set consists of ordered pairs. The term "directed graph" is often abbreviated as* ***digraph****.*

Although graphs are defined abstractly as above, it's very common to draw *diagrams* to represent them. These are drawings in which the vertices are shown as points or disks and the edges as line segments or arcs. Figure 1.4 illustrates the graphical convention used to mark the distinction between directed and undirected edges: the former are drawn as line segments or arcs, while the latter are shown as arrows. A directed edge $e = (a, b)$ appears as an arrow that points from a to b.

Sometimes one sees graphs with more than one edge[3] connecting the same two vertices; the Königsberg Bridge graph is an example. Such edges are called *multiple* or *parallel* edges. Additionally, one sometimes sees graphs with edges of the form $e = (v, v)$. These edges, which connect a vertex to itself, are called *loops* or *self loops*. All these terms are illustrated in Figure 1.5.

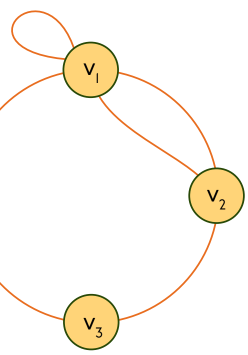

Figure 1.5: *A graph whose edge set includes the self loop (v_1, v_1) and two parallel copies of the edge (v_1, v_2).*

It is important to bear in mind that diagrams such as those in Figures 1.3–1.5 are only illustrations of the edges and vertices. In particular, the arcs representing edges may cross, but this does not necessarily imply anything: see Figure 1.6.

Remark. In this course when we say "graph" we will normally mean an undirected graph that contains no loops or parallel edges: if you look in other books you may

[3] In this case it is a slight abuse of terminology to talk about the edge "set" of the graph, as sets contain only a single copy of each of their elements. Very scrupulous books (and students) might prefer to use the term *edge list* in this context, but I will not insist on this nicety.

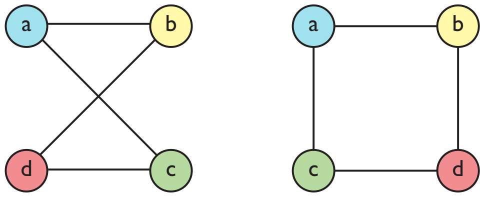

Figure 1.6: *Two diagrams for the same graph: the crossed edges in the leftmost version do not signify anything.*

see such objects referred to as *simple graphs*. By contrast, we will refer to a graph that contains parallel edges as a *multigraph*.

Definition 1.4. *Two vertices $a \neq b$ in an undirected graph $G(V, E)$ are said to be **adjacent** or to be **neighbours** if $(a, b) \in E$. In this case we also say that the edge $e = (a, b)$ is **incident on** the vertices a and b.*

Definition 1.5. *If the directed edge $e = (u, v)$ is present in a directed graph $H(V', E')$ we will say that u is a **predecessor** of v and that v is a **successor** of u. We will also say that u is the **tail** or **tail vertex** of the edge (u, v), while v is the **tip** or **tip vertex**.*

1.3 Standard examples

In this section I'll introduce a few families of graphs that we will refer to throughout the rest of the term.

The complete graphs K_n

The *complete graph* K_n is the undirected graph on n vertices whose edge set includes every possible edge. If one numbers the vertices consecutively the edge and vertex set are

$$\begin{aligned} V &= \{v_1, v_2, \ldots, v_n\} \\ E &= \{(v_j, v_k) \,|\, 1 \leq j \leq (n-1),\ (j+1) \leq k \leq n\}. \end{aligned}$$

There are thus

$$|E| = \binom{n}{2} = \frac{n(n-1)}{2}$$

edges in total: see Figure 1.7 for the first few examples.

The path graphs P_n

These graphs are formed by stringing n vertices together in a path. The word "path" actually has a technical meaning in graph theory, but you needn't worry about that

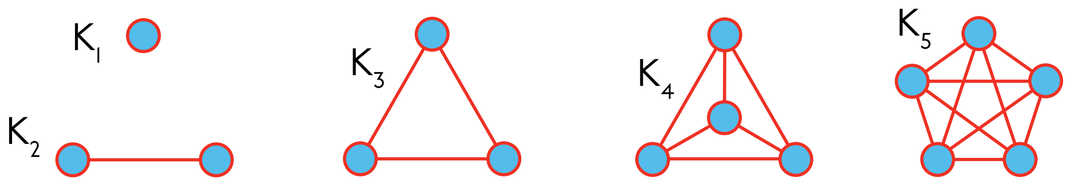

Figure 1.7: *The first five members of the family K_n of complete graphs.*

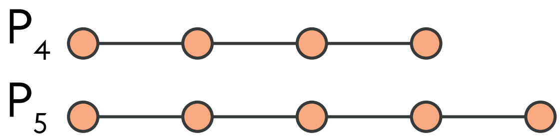

Figure 1.8: *Diagrams for the path graphs P_4 and P_5.*

today. P_n has vertex and edge sets as listed below,

$$\begin{aligned} V &= \{v_1,\, v_2,\, \ldots,\, v_n\} \\ E &= \{(v_j, v_{j+1}) \,|\, 1 \leq j < n\}\,, \end{aligned}$$

and Figure 1.8 shows two examples.

The cycle graphs C_n

The *cycle graph* C_n, sometimes also called the *circuit graph*, is a graph in which $n \geq 3$ vertices are arranged in a ring. If one numbers the vertices consecutively the edge and vertex set are

$$\begin{aligned} V &= \{v_1,\, v_2,\, \ldots,\, v_n\} \\ E &= \{(v_1, v_2),\, (v_2, v_3),\, \ldots,\, (v_j, v_{j+1}),\, \ldots,\, (v_{n-1}, v_n),\, (v_n, v_1)\}\,. \end{aligned}$$

C_n has n edges that are often written (v_j, v_{j+1}), where the subscripts are taken to be defined periodically so that, for example, $v_{n+1} \equiv v_1$. See Figure 1.9 for examples.

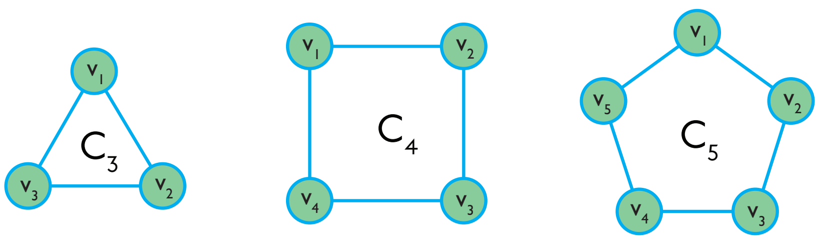

Figure 1.9: *The first three members of the family C_n of cycle graphs.*

The complete bipartite graphs $K_{m,n}$

The *complete bipartite graph* $K_{m,n}$ is a graph whose vertex set is the union of a set V_1 of m vertices with second set V_2 of n different vertices and whose edge set includes every possible edge running between these two subsets:

$$\begin{aligned} V &= V_1 \cup V_2 \\ &= \{u_1, \ldots, u_m\} \cup \{v_1, \ldots, v_n\} \\ E &= \{(u, v) \mid u \in V_1, \, v \in V_2\} \,. \end{aligned}$$

$K_{m,n}$ thus has $|E| = mn$ edges: see Figure 1.10 for examples.

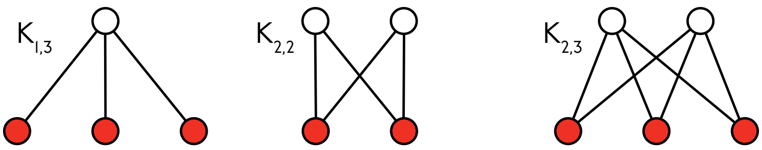

Figure 1.10: *A few members of the family $K_{m,n}$ of complete bipartite graphs. Here the two subsets of the vertex set are illustrated with colour: the white vertices constitute V_1, while the red ones form V_2.*

There are other sorts of bipartite graphs too:

Definition 1.6. *A graph $G(V, E)$ is said to be a* ***bipartite graph*** *if*

- *it has a nonempty edge set: $E \neq \emptyset$ and*
- *its vertex set V can be decomposed into two nonempty, disjoint subsets*

  $$V = V_1 \cup V_2 \ \text{ with } \ V_1 \cap V_2 = \emptyset \ \text{ and } \ V_1 \neq \emptyset \ \text{ and } \ V_2 \neq \emptyset$$

  in such a way that all the edges $(u, v) \in E$ contain a member of V_1 and a member of V_2.

The cube graphs I_d

These graphs are specified in a way that's closer to the purely combinatorial, set-theoretic definition of a graph given above. I_d, the *d-dimensional cube graph*, has vertices that are strings of d zeroes or ones, and all possible labels occur. Edges connect those vertices whose labels differ in exactly one position. Thus, for example, I_2 has vertex and edge sets

$$V = \{00, \, 01, \, 10, \, 11\} \qquad \text{and} \qquad E = \{(00, 01), \, (00, 10), \, (01, 11), \, (10, 11)\} \,.$$

Figure 1.11 shows diagrams for the first few cube graphs and these go a long way toward explaining the name. More generally, I_d has vertex and edge sets given by

$$\begin{aligned} V &= \left\{w \mid w \in \{0, 1\}^d\right\} \\ E &= \{(w, w') \mid w \text{ and } w' \text{ differ in a single position}\} \,. \end{aligned}$$

This means that I_d has $|V| = 2^d$ vertices, but it's a bit harder to count the edges. In the last part of today's lecture we'll prove a theorem that enables one to show that I_d has $|E| = d \, 2^{d-1}$ edges.

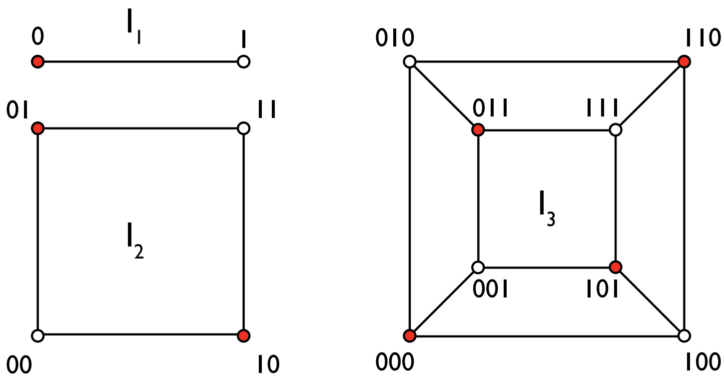

Figure 1.11: *The first three members of the family I_d of cube graphs. Notice that all the cube graphs are bipartite (the red and white vertices are the two disjoint subsets from Definition 1.6), but that, for example, I_3 is* not *a complete bipartite graph.*

1.4 A first theorem about graphs

I find it wearisome to give, or learn, one damn definition after another and so I'd like to conclude the lecture with a small, but useful theorem. To do this we need one more definition:

Definition 1.7. *In an undirected graph $G(V, E)$ the* ***degree*** *of a vertex $v \in V$ is the number of edges that include the vertex. One writes* $\deg(v)$ *for "the degree of v".*

So, for example, every vertex in the complete graph K_n has degree $n-1$, while every vertex in a cycle graph C_n has degree 2; Figure 1.12 provides more examples. The generalization of degree to directed graphs is slightly more involved. A vertex v in a digraph has two degrees: an *in-degree* that counts the number of edges having v at their tip and an *out-degree* that counts number of edges having v at their tail. See Figure 1.13 for an example.

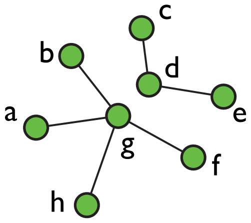

v	a	b	c	d	e	f	g	h
$\deg(v)$	1	1	1	2	1	1	4	1

Figure 1.12: *The degrees of the vertices in a small graph. Note that the graph consists of two "pieces".*

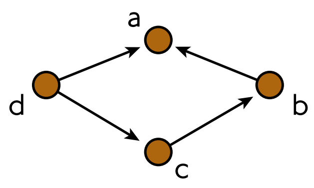

v	$\deg_{in}(v)$	$\deg_{out}(v)$
a	2	0
b	1	1
c	1	1
d	0	2

Figure 1.13: *The degrees of the vertices in a small digraph.*

Once we have the notion of degree, we can formulate our first theorem:

Theorem 1.8 (Handshaking Lemma, Euler 1736)**.** *If $G(V, E)$ is an undirected graph then*

$$\sum_{v \in V} \deg(v) = 2|E|. \tag{1.1}$$

Proof. Each edge contributes twice to the sum of degrees, once for each of the two vertices on which it is incident. □

The following two results are immediate consequences:

Corollary 1.9. *In an undirected graph there must be an even number of vertices that have odd degree.*

Corollary 1.10. *The cube graph I_d has $|E| = d\,2^{d-1}$.*

The first is fairly obvious: the right hand side of (1.1) is clearly an even number, so the sum of degrees appearing on the left must be even as well. To get the formula for the number of edges in I_d, note that it has 2^d vertices, each of degree d, so the Handshaking Lemma tells us that

$$2|E| \;=\; \sum_{v \in V} \deg(v) \;=\; 2^d \times d$$

and thus $|E| = (d \times 2^d)/2 = d\,2^{d-1}$.