

Introduction

A directed graph is a combinatorial object consisting of vertices and oriented edges joining pairs of vertices. We can represent such a graph by operators on a Hilbert space \mathcal{H} : the vertices are represented by mutually orthogonal closed subspaces, or more precisely the projections onto these subspaces, and the edges by operators between the appropriate subspaces. The graph algebra is, loosely speaking, the C^* -subalgebra of $B(\mathcal{H})$ generated by these operators.

When the graph is finite and highly connected, the graph algebras coincide with a family of C^* -algebras first studied by Cuntz and Krieger in 1980 [16]. The Cuntz-Krieger algebras were quickly recognised to be a rich supply of examples for operator algebraists, and also cropped up in some unexpected situations [87, 131]. In the past ten years there has been a great deal of interest in graph algebras associated to infinite graphs, and these have arisen in new contexts: in non-abelian duality [83, 24], as deformations of commutative algebras [56, 57, 58], in non-commutative geometry [12], and as models for the classification of simple C^* -algebras [136].

Graph algebras have an attractive structure theory in which algebraic properties of the algebra are related to combinatorial properties of paths in the directed graph. The fundamental theorems of the subject are analogues of those proved by Cuntz and Krieger, and include a uniqueness theorem and a description of the ideals in graph algebras. But we now know much more: just about any C^* -algebraic property a graph algebra might have can be determined by looking at the underlying graph.

Our goals here are to describe the structure theory of graph algebras, and to discuss two particularly promising extensions of that theory involving the topological graphs of Katsura [73] and the higher-rank graphs of Kumjian and Pask [81]. We provide full proofs of the fundamental theorems, and also when we think some insight can be gained by proving a special case of a published result or by taking an alternative route to it. Otherwise we concentrate on describing the main ideas and giving references to the literature.

Outline. The core material is in the first four chapters, where we discuss the uniqueness theorems and the ideal structure. These theorems were first proved for infinite graphs by realising the graph algebra as the C^* -algebra of a locally compact groupoid, and applying results of Renault [83, 82]. There are now several other approaches to this material; the elementary methods we use here are based on the original arguments of Cuntz and Krieger, but incorporate several simplifications which have been made over the years. These techniques work best for the row-finite graphs in which each vertex receives just finitely many edges; in Chapter 5 we describe a method of Drinen and Tomforde for reducing problems to the row-finite case.

In Chapter 6, we describe how graph algebras provided important insight into some problems in non-abelian duality for crossed products of C^* -algebras, and then in Chapter 7 we calculate the K -theory of graph algebras. In Chapter 8, we give an introduction to correspondences and Cuntz-Pimsner algebras, using graph algebras as motivation for the various constructions. This is important material for researchers in many areas: many interesting C^* -algebras are by definition the Cuntz-Pimsner algebra of some particular correspondence (see [68, 73, 92], for example), and the general theory seems to be a powerful tool. This is certainly the case for the C^* -algebras of topological graphs, which are the subject of Chapter 9. In the last chapter, we discuss higher-rank graphs.

Chapters 5, 6, 7 and 10 are essentially independent of each other, and can be read in any order after the first four chapters. Chapter 9, on the other hand, requires familiarity with Chapter 8.

Conventions. We use the standard conventions of our subject. Thus, for example, homomorphisms between C^* -algebras are always $*$ -preserving, and representations of C^* -algebras are homomorphisms into the algebra $B(\mathcal{H})$ of bounded linear operators on a Hilbert space \mathcal{H} ; ideals are always assumed to be closed and two-sided. As a more personal convention, we always denote the inner product of two elements h, k of a Hilbert space by $(h | k)$, and reserve the notation $\langle x, y \rangle$ for C^* -algebra-valued inner products, which for us are always conjugate-linear in the first variable (see Chapter 8).

There is, unfortunately, no standard set of conventions for graph algebras: one can choose to have the partial isometries representing edges going either in the same direction as the edge or in the opposite direction. Both conventions are used in the literature, and both have their advantages and disadvantages. We have opted to have the partial isometries going in the same direction as the edges. The main disadvantage of this is that we have to adopt a rather unnatural notation for paths in graphs (see Remark 1.13). However, when edges represent morphisms in a category, as they do in higher-rank graphs, for example, this convention means that composition of morphisms is compatible with multiplication of operators in $B(\mathcal{H})$. After several years of having bright students who talk about higher-rank graphs all the time, I have found it easier to adopt their conventions. (There was no chance they were going to change. . .) Nevertheless, I sympathise with those to whom changing seems an unnecessary nuisance, and I apologise to them.

Background. While writing these notes, I have been addressing a reader who has taken a first course in C^* -algebras, covering the Gelfand-Naimark Theorems, the continuous functional calculus, and positivity. This material is covered in some form or other in most of the standard books on the subject, and the first 3 chapters of [93], for example, contain everything we need.

For those with slightly different backgrounds, it might be helpful to mention some non-trivial facts about C^* -algebras which we use frequently.

- (1) Every homomorphism between C^* -algebras is norm-decreasing, and every injective homomorphism is norm-preserving. In particular, every automorphism and every isomorphism of C^* -algebras preserves the norm.
- (2) The range $\phi(A)$ of every homomorphism $\phi : A \rightarrow B$ between C^* -algebras is closed, and is therefore a C^* -subalgebra of B .

- (3) Every C^* -algebra has a faithful representation as a C^* -algebra of bounded operators on Hilbert space. This is often used when we want to apply results about representations of C^* -algebras to more general homomorphisms.

Points (1) and (2) give the theory of C^* -algebras a rather algebraic flavour. However, it is essential in (1) and (2) that the algebras are complete and that the homomorphisms are everywhere defined. We often work with maps defined on dense $*$ -subalgebras of C^* -algebras, and then we have to establish norm estimates to be sure that the maps extend to the C^* -algebras.

We have included an appendix in which we discuss some material which might not be covered in a first course, and which might be hard to locate in the literature. First we show how geometric properties of projections and partial isometries on a Hilbert space \mathcal{H} can be encoded using the $*$ -algebraic structure of $B(\mathcal{H})$. Then we look at some standard tricks which allow us to identify C^* -algebras as matrix algebras or direct sums and direct limits of such algebras.