Cambridge University Press
978-0-521-49797-8 - Surveys in Combinatorics, 1995
Edited by Peter Rowlinson
Excerpt
More information

EUCLIDEAN GEOMETRY OF DISTANCE REGULAR GRAPHS

C. D. Godsil [1]

Combinatorics and Optimization
University of Waterloo
Waterloo, Ontario
Canada N2L 3G1

ABSTRACT

A graph is distance regular if it is connected and, given any two vertices u and v at distance i, the number of vertices x at distance j from u and k from v is determined by the triple (i, j, k). Distance regular graphs are interesting because of their connections with coding theory, design theory and finite geometry.

We can introduce geometric methods into the study of distance regular graphs as follows. Let X be a graph on n vertices and identify the i-th vertex of X with the i-th standard basis vector e_i in $\mathbf{R}^n$. Let θ be an eigenvalue of (the adjacency matrix of) X and let U be the corresponding eigenspace. We then associate to the i-th vertex of X the image of e_i under orthogonal projection onto U. If U has dimension m then we have a mapping from $V(X)$ into $\mathbf{R}^n$. If X is distance regular, then it can be shown that the image of $V(X)$ lies in a sphere centred at the origin, and that the cosine of the angle between the vectors representing two vertices u and v is determined by the distance between them in X. In this paper we survey some of the applications of these methods.

---

[1] Support from grant OGP0093041 of the National Sciences and Engineering Council of Canada is gratefully acknowledged.

1. Introduction

We define a *representation* of a graph G in $\mathbf{R}^m$ to be a map, ρ, from $V(G)$ into $\mathbf{R}^m$, such that for any two vertices u and v, the inner product

$$\langle \rho(u), \rho(v) \rangle$$

is determined by the distance $\operatorname{dist}(u, v)$ between u and v in G. Because any vertex is at distance 0 from itself, this implies that the image of $V(G)$ under ρ lies on a sphere centred at the origin. We will always assume that ρ is not the zero map, although it might map all vertices of G to the same non-zero vector.

To obtain a non-trivial example, consider the dodecahedron embedded as a regular polytope in $\mathbf{R}^3$ with centre of mass at the origin. We may assume that its vertices lie on the unit sphere. If G is the 1-skeleton of this polytope then the map that assigns to each vertex of G the vector representing it in $\mathbf{R}^3$ is a representation. Further examples can be obtained from the other Platonic solids.

Representations provide us with the opportunity to apply geometric methods to graph theory, in particular to the study of distance-regular graphs.

2. Representations

We first introduce one concept from linear algebra. The *Gram matrix* of a set of vectors $v_1, \ldots, v_n$ is the $n \times n$ matrix M with $M_{i,j}$ equal to $\langle v_i, v_j \rangle$. It is not hard to see that a Gram matrix is positive semi-definite and symmetric. The converse is also true, any symmetric positive semi-definite matrix is a Gram matrix.

If G is a graph with diameter d, let G_i denote the graph with the same vertex set as G, with two vertices adjacent in G_i if and only if they are at distance i in G. Let A_i denote the adjacency matrix of G_i, with the understanding that $A_0 = I$. We call $A_0, \ldots, A_d$ the *distance matrices* of G. If ρ is a representation of G in $\mathbf{R}^m$, then the matrix $M(\rho)$ has rows and columns indexed by the vertices of G and, if $u, v \in V(G)$, then

$$(M(\rho))_{u,v} := \langle \rho(u), \rho(v) \rangle.$$

The matrix $M(\rho)$ is a linear combination of the matrices A_i. It is a Gram matrix, of the vectors ρ for u in $B(G)$, and so it is positive semi-definite. As every symmetric positive semi-definite matrix is a Gram matrix, we see conversely that each positive semi-definite matrix in the span of $A_0, \dots, A_d$ determines a representation of G.

A subgraph H of G is *isometric* if $\text{dist}_H(u, v) = \text{dist}_G(u, v)$ for any two vertices u and v of H. It follows that if ρ is a representation of G then its restriction to any isometric subgraph H is a representation of H.

We have not required that a representation be injective. We can sometimes show that if $\text{dist}(u, v) \leq r$, then $\rho(u) \neq \rho(v)$; if this is the case we will say that ρ is r-*injective*.

By way of introduction, we offer the following.

Lemma 2.1. *There are only finitely many graphs G such that both G and its complement admit an injective representation into $\mathbf{R}^m$.*

Proof. If ρ is injective, then the image of any clique is a regular simplex. Hence if G has a representation in $\mathbf{R}^m$ then there are no cliques in G with more than $m + 1$ vertices in them. Because the complement $\overline{G}$ of G also has a representation on $\mathbf{R}^m$, there can be no independent set in G with size greater than $m + 1$. Our claim follows by Ramsey's theorem. □

Lemma 2.2. *If G admits a 2-injective representation into $\mathbf{R}^m$, then the maximum valency of a vertex in G is bounded by a function of m.*

Proof. If N is the neighbourhood of a vertex in G, then two distinct vertices in N are at distance 1 or 2 in G. Hence if ρ is a 2-injective representation of G into $\mathbf{R}^m$, then the restriction of ρ to N is a representation of both N and $\overline{N}$ in an $(m - 1)$-dimensional affine subspace of $\mathbf{R}^m$. So, by the previous lemma, $|V(N)|$ is bounded by a function of m. □

Cambridge University Press
978-0-521-49797-8 - Surveys in Combinatorics, 1995
Edited by Peter Rowlinson
Excerpt
More information

If we compute an explicit expression for the function in this lemma using Ramsey theory, we obtain an exponential bound for the maximum valency of G. We will see in the next section that this is far too large.

3. Geometry

As the image of a graph under a representation is a set of points in $\mathbf{R}^m$, we need some terminology for such subsets. Let S be a subset of the unit sphere in $\mathbf{R}^m$. The *degree* of S is the size of its *degree set*

$$\{\langle x,y\rangle : x,y \in S,\ x \neq y\}.$$

A set with degree s is called an *s-distance set*. If f is a polynomial in m variables then f is a function on S. We say that S is a *spherical t-design* if the average over the points of S of any polynomial f with degree at most t is equal to the average value of f over the unit sphere. This terminology is based on an analogy with the theory of t-designs; roughly speaking a spherical t-design is a finite approximation to the unit sphere, whereas a t-(v,k,λ) design is an approximation to the set of all k-subsets of a fixed set of v elements. For more on this viewpoint, see [7: Chapter 14]. For information on spherical designs see [4] and [7: Chapter 13]. The largest integer t such that S is a t-design is the *strength* of S. A subset S of the unit sphere is a 1-design if and only if the sum of the vectors in it is the zero vector.

The following bound is a combination of results from [4], in particular Theorems 2.4 and 5.11.

Theorem 3.1. *Let S be a subset of the unit sphere in $\mathbf{R}^m$ with degree s and strength t, and define $f(r)$ to be $\binom{m+r-1}{r}+\binom{m+r-2}{r-1}$. Then*

$$f(\lfloor t/2\rfloor) \leq |S| \leq f(s);$$

further, if one of these inequalities is an equality, then so is the other.

A subset S for which equality holds in the bound of this theorem is called a *tight* spherical design. It implies that a 2-distance set in $\mathbf{R}^m$ has size at most $m(m+3)/2$, and therefore this value may be taken as the bound of Lemma 2.2.

CAMBRIDGE

Cambridge University Press
978-0-521-49797-8 - Surveys in Combinatorics, 1995
Edited by Peter Rowlinson
Excerpt
More information

Theorem 3.2. *Let S be a subset of the unit sphere in $\mathbf{R}^m$ with degree s and strength t. If $x \in S$ and α belongs to the degree set of S then $\{y \in S : \langle x, y\rangle = \alpha\}$ is a $(t+1-s)$-design.*

Let S be a finite set of points on the unit sphere in $\mathbf{R}^m$. If α lies in the degree set of S, let A_α be the $(0,1)$-matrix with rows and columns indexed by S, and with $(A_\alpha)_{u,v} = 1$ if and only if $\langle u, v\rangle = \alpha$.

Theorem 3.3. *Let S be a subset of the unit sphere in $\mathbf{R}^m$ with degree s and strength t, and let Δ be its degree set. If $t \geq 2s-2$, then the matrices A_α for α in Δ, together with the identity matrix, form an s-class association scheme.*

A finite set of symmetric $(0,1)$-matrices $A_0, \ldots, A_d$ forms an *association scheme* if:

(a) $A_0 = I$,
(b) $\sum_\alpha A_\alpha = J$,
(c) for all i and j, the product A_iA_j is a linear combination of $A_0, \ldots, A_d$.

For more on association schemes see, e.g., [3, 7]. It is not too hard to show that $t \leq 2s$ for any finite subset of the unit sphere, so the lower bound on t in this theorem is quite strong. An important consequence of the last theorem is that a 2-distance set with strength at least 2 gives rise to a complementary pair of strongly regular graphs. Further, all strongly regular graphs can be obtained from such subsets of the unit sphere.

4. Distance Regular Graphs

A graph G with diameter d is *distance regular* if, for $i = 1, \ldots, d$, the distance matrix A_i is a polynomial in A_1 with degree i. In more combinatorial terms, G has the property that if integers i, j and k are given and u and v are vertices in G at distance i then the number of vertices in G at distance j from u and distance k from v is independent of the choice of u and v. A graph is *distance transitive* if, given any two ordered pairs of vertices (u, v) and (u', v') such that $\text{dist}(u, v) = \text{dist}(u', v')$, there is an automorphism of

G that maps (u, v) to (u', v'). It is easy to see that a distance transitive graph is distance regular. The 1-skeletons of the Platonic solids in $\mathbf{R}^3$ are all distance-transitive; this is an almost immediate consequence of the fact that they are regular polytopes. For a complete introduction to the theory of distance regular (and distance transitive) graphs, see [3].

We point out that there are vast numbers of distance regular graphs that have trivial automorphism group; distance regularity does not imply the existence of any non-identity automorphisms in general.

The *Johnson graph* $J(v, \ell)$ has all ℓ-subsets of a fixed set V of size v as its vertices, with two ℓ-subsets adjacent if they intersect in exactly $\ell - 1$ points. This is a distance transitive graph, and therefore distance regular. Let Q be a fixed set of size q. The graph $J(v, 2)$ is also known as the line graph of the complete graph. The *Hamming graph* $H(n, q)$ has vertex set Q^n, and two elements of Q^n are adjacent if they differ in exactly one coordinate. (The set Q is often taken to be a finite field, but we do not even require it to have prime power order.) The Hamming graph $H(n, 2)$ is n-cube. The Hamming graphs are all distance transitive.

The distance matrices of a distance regular graph with diameter d form a basis of a commutative algebra of dimension $d + 1$ over the reals, called the *Bose-Mesner algebra*. Thus, if ρ is a representation of G, then $M(\rho)$ is a positive semi-definite matrix in the Bose-Mesner algebra of G; conversely every such matrix gives rise to a representation of G.

One class of positive semi-definite matrices can be obtained using projections. Let θ be an eigenvalue of $A = A_1$ with multiplicity m, let U be the corresponding eigenspace and let E_θ be the matrix representing orthogonal projection on U. Then E_θ can be shown to be a polynomial in A, and therefore it lies in the Bose-Mesner algebra of G. Hence we obtain a representation of G in $\mathbf{R}^m$, which we will call an *eigenspace representation* of G.

As E_θ lies in the span of the matrices $A_0, \ldots, A_d$, its diagonal entries are all equal. Because $E_\theta^2 = E_\theta$, all eigenvalues of E_θ are equal to 0 or 1; therefore $\operatorname{tr} E_\theta = \operatorname{rk} E_\theta$ and each diagonal entry of E_θ is equal to $m/|V(G)|$.

So there are constants $w_0, \dots, w_d$, with $w_0 = 1$, such that

$$\frac{|V(G)|}{m} E_\theta = \sum_{i=0}^{d} w_i A_i. \tag{4.1}$$

As this matrix is positive semi-definite, each principal 2×2 submatrix has non-negative determinant, which implies that

$$|w_i| \le 1.$$

We can view w_i as the cosine of the angle between $\rho(u)$ and $\rho(v)$, for any two vertices u and v at distance i in G. We call $w_0, \dots, w_d$ the *sequence of cosines* of G; this sequence depends on the eigenvalue θ. We can summarise our conclusions as follows.

Lemma 4.1. *Let G be a distance regular graph with diameter d, let θ be an eigenvalue of G, with cosine sequence $w_0, \dots, w_d$ and let ρ be the corresponding representation of G. Then*

$$M(\rho) = \sum_{i=0}^{d} w_i A_i.$$ □

If $A = A_1$ then $AM(\rho) = \theta M(\rho)$. As A is a $(0,1)$-matrix, this implies that

$$\theta \rho(u) = \sum_{v \sim u} \rho(v). \tag{4.2}$$

(We write $v \sim u$ to denote that v is adjacent to u.) We also have the following.

Lemma 4.2. *Let G be a distance regular graph with valency k and let θ be a non-trivial eigenvalue of G. If ρ is the representation of G on the eigenspace belonging to θ, then $\rho(G)$ is a spherical 2-design.* □

If G is a distance regular graph with valency k, then k is an eigenvalue of G. The associated eigenspace is spanned by $\mathbf{1}$, the vector with all entries equal to 1, and hence the representation we obtain is not interesting. Further, $-k$ is an eigenvalue of G if and only if G is bipartite; the associated eigenspace is, once again, 1-dimensional. We will call an eigenvalue θ of G *non-trivial* if $\theta \neq \pm k$.

5. Cosines

If G is distance regular, then there are constants a_i, b_i and c_i such that

$$AA_i = b_{i-1}A_{i-1} + a_iA_i + c_{i+1}A_{i+1}. \tag{5.1}$$

(This is essentially the matrix formulation of the combinatorial definition of distance regularity given at the start of the previous section.) Here $a_0 = c_0 = 0$, $c_1 = 1$ and $b_d = 0$. We denote the valency of G by k, so $b_0 = k$. Further, a_1 is the number of triangles on an edge of G and c_2 is the number of common neighbours of two vertices at distance 2 in G. By way of example, for the Johnson graph $J(v, k)$ we have

$$b_i = (k - i)(v - k - i), \quad a_i = i(v - 2i), \quad c_i = i^2,$$

when $i = 0, \dots, k$. In all cases we have that

$$b_i + a_i + c_i = k, \qquad i = 0, \dots, d.$$

If x is at distance i from u in the distance regular graph G then there are b_i neighbours of u at distance $i+1$ from x, a_i at distance i and c_i at distance $i - 1$. Hence, if ρ is the representation of G on the eigenspace belonging to θ and we take the inner product of both sides of (4.2) with $\rho(x)$, then we obtain:

$$\theta w_i = c_iw_{i-1} + a_iw_i + b_iw_{i+1}. \tag{5.2}$$

As $w_0 = 1$, this implies that $\theta = b_0w_1$ and thus:

$$w_1 = \frac{\theta}{k}. \tag{5.3}$$

Further, $\theta w_1 = 1 + a_1w_1 + b_1w_2$ and $1 + a_1 + b_1 = k$, whence

$$w_2 = \frac{(\theta - a_1)w_1 - 1}{b_1} = \frac{\theta^2 - a_1\theta - k}{k(k - 1 - a_1)}. \tag{5.4}$$

We see that if $w_1 = 1$, then $\theta = k$ and the corresponding eigenspace is spanned by the vector **1**. If $w_1 = -1$ then $\theta = -k$. In this case G must be bipartite, the multiplicity of θ is 1 and the corresponding eigenvector takes value 1 on one colour class of G and -1 on the other. The next result answers most questions about injectivity.

Lemma 5.1. *Let G be a connected distance regular graph with diameter d and valency k, where $k \geq 3$. Let $\theta_0 > \theta_1 > \cdots > \theta_d$ be the eigenvalues of G. The eigenspace representation corresponding to θ_i is not injective if and only if*

(a) $i = 0$, *(that is,* $\theta = k$*), or*

(b) $\theta = -k$ *and G is bipartite, or*

(c) i *is even and G is antipodal.* □

A distance regular graph G is said to be *antipodal* if G_d is not connected, in which case the components of G_d all have the same size. Complete multipartite graphs are precisely the antipodal distance regular graphs with diameter 2. The proof of Lemma 5.1 given on [7: p. 265] also yields the following.

Lemma 5.2. *Let G be a distance regular graph with valency k, and let θ be a non-trivial eigenvalue of G. If G is not complete multipartite then the representation belonging to θ is 2-injective.* □

The number of *sign-changes* in a sequence $a_0, \ldots, a_d$ is the number of indices i such that $a_i a_{i+1} < 0$. Deleting the terms of a sequence that are equal to 0 does not change the number of sign-changes in it. A proof of the next result appears in [7: Lemma 13.3.1].

Lemma 5.3. *Let G be a distance regular graph with diameter d and let θ be the i-th largest eigenvalue of G. Then the cosine sequence $w_0, \ldots, w_d$ has exactly i sign-changes, and the sequence $w_0 - w_1, \ldots, w_{d-1} - w_d$ has exactly $i - 1$ sign-changes.* □

From (5.4) above we see that

$$1 - w_2 = \frac{k-\theta}{k} \frac{\theta - a_1 + k}{k - a_1 - 1}. \tag{5.5}$$

As $|w_2| \leq 1$, one consequence of this is that

$$a_1 - k \leq \theta,$$

for any eigenvalue θ of G. We also have

$$1 - 2w_1 + w_2 = \frac{k-\theta}{k}\,\frac{k-a_1-2-\theta}{k-a_1-1}. \qquad (5.6)$$

The parameter $1-2w_1+w_2$ turns out to be important; partial evidence for this is provided by the next result, taken from [3: Prop. 4.4.9].

Theorem 5.4. *Let G be a distance regular graph with diameter d and valency k, let θ be a non-trivial eigenvalue of G and let $w_0, \dots, w_d$ be the corresponding sequence of cosines. If G contains an induced copy of C_4, then $1 - 2w_1 + w_2 \geq 0$. If equality holds, then $\theta = k - a_1 - 2$ and $w_i = 1 - i(a_1+2)/k$.*

Proof. Let ρ be the representation of G on the eigenspace belonging to θ and let u, v, x and y be the images under ρ of the vertices in an induced 4-cycle in G; where u and x represent vertices at distance 2 in G, as do v and y. Then

$$0 \leq \|u+x-v-y\|^2 = 4(1-2w_1+w_2) \qquad (5.7)$$

proving our claim.

Suppose equality holds in (5.7), and suppose that z is the image under ρ of a vertex at distance $i-1$ from u and distance $i+1$ from x. Then $u+x-v-y=0$ and therefore

$$0 = \langle z, u+x-v-y\rangle = w_{i-1} + w_{i+1} - 2w_i.$$

Given this it is not hard to show that $w_i = 1 - i\frac{a_1+2}{k}$. □

Lemma (Delsarte) 5.5. *Let G be a distance regular graph with diameter d and valency k. If C is a clique in G of size c, then $c \leq 1 - (k/\theta_d)$.*

Proof. If θ is an eigenvalue of G, then the $c \times c$ matrix

$$I + w_1(J-I)$$

is positive semi-definite, because it is a positive constant times a principal submatrix of the positive semi-definite matrix E_θ. Hence all its eigenvalues