This report is submitted for credit for MATHS 783, semester II 2004.

Name: Robin Christian
Student ID: 2376183
Supervisor: Dr Paul Bonnington

Algebraic Graph Theory

Chris Godsil & Gordon Royle

During the course of this semester I have read chapters (1), (2), 3, 4, 5, (8), 12, 13 and 15 of Algebraic Graph Theory by Chris Godsil and Gordon Royle. The chapters in brackets were revision or introductory material. Briefly, the content of each (important) chapter was:

Chapter 3: Properties of vertex-transitive and edge-transitive graphs, connectivity of transitive graphs, matchings, Hamilton paths & cycles.

Chapter 4: Arc-transitive and distance-transitive graphs, s-arc regularity, the Coxeter graph and Tutte's 8-cage.

Chapter 5: Partial linear spaces, generalized polygons, Moore graphs, the Hoffman-Singleton graph.

Chapter 12: Generalized line graphs, the characterization of all graphs with minimum eigenvalue at least -2 using root systems.

Chapter 13: The Laplacian of a graph, counting spanning trees, representations, results on the second-lowest eigenvalue of the Laplacian, interlacing, conductance.

Chapter 15: Matroids and their relationship with graphs and codes, the rank polynomial, deletion-contraction.

This report consists of hints for selected exercises from each of these chapters. Most of the exercises I completed are included here, although I left out "semi-proved" exercises, and a few (most of these were from Chapter 13) where I had a complete proof but I could not find a way to hint at or summarize the solution concisely. Scribbled versions of these proofs are in the notes handed in with this report. I have indicated some of my favourite exercises with a *.

In a few places my notation differs from that of the text. I have used x·y for the inner product of x and y instead of <x,y>, radians instead of degrees, A

+B for the union of sets A and B, and A * B for the Cartesian product of graphs A and B.

Hints on Selected Exercises

Chapter 3

1. Label the vertices of the original K_5 in the construction {1, 2, 3, 4, 5} clockwise, and the 'duplicate' vertex i' for each i. Label the "edge" vertex connected to i, j, i' and j' as ij. Then label the vertices in Figure 3.2 {14, 25, 13, 24, 35, 45, 34, 23, 12, 15} down the main diagonal and {1, 2, 3, 4, 5, 5', 4', 3', 2', 1'} down the secondary diagonal.

It should be clear that this provides an isomorphism, and also that any permutation of {1, 2, 3, 4, 5} acting on vertex labels is an automorphism. These automorphisms are transitive on the edges, but no automorphism can map e.g. 1 to 12, because it would also have to map 1' to 12.

2. There are only two groups on ten vertices, Z_{10} and D_5. Imagine that the Petersen graph is X(G, C) where G is either of these groups. Considering the possibilities for C it is quickly seen that if X(G, C) is cubic it must contain a 4-cycle.

3. The dodecahedron is a 2-fold cover of the Petersen graph (this can be seen by identifying opposite points on the dodecahedron). If the dodecahedron were a Cayley graph X(G, C) and f the homomorphism from the dodecahedron to the Petersen graph then the Petersen graph would be the Cayley graph X(f(G), f(C)). It is apparent that each element of C would have a distinct image under f by considering the neighbors of 1.

4.If X(G, C) is connected then for any g in G there is a path from e to g^{-1}. Multiplying along this path shows that g is a product of elements of C. Conversely if $g = c_1 \ldots c_n$, then we have the following path in X(G, C): $\{e, c_n, c_{n-1}c_n, \ldots, c_1 \ldots c_n = g\}$.

5. All edges must be between the cosets A_n and (12) A_n.

6. If G is a vertex-transitive graphs on p vertices, then $|x^{Aut(G)}| = p$ for each vertex x. Thus |Aut(G)| divides p by the orbit-stabiliser theorem. By Frobenius's Lemma, there is g in Aut(G) with order p. Show that <g> acts regularly on G, then G is a Cayley graph by Lemma 3.7.2.

7(*). No. Consider $\Sigma|S\cap S^g|$. Since G is transitive on V, $|G| = |G_x||x^G|$ so $|G_x| = |G|/|V|$. Since G is transitive there is a 1-1 correspondence between G_x and the elements of G that map x to y. Deduce that $\Sigma|S\cap S^g| = |S|^2|G|/|V|$. Since $|S\cap S^g| \geq c$ for $g \neq e$, and $|S\cap S^e| = |S|$, obtain $\Sigma|S\cap S^g| > c|G|$ and the result follows.

8. Suppose that G is Abelian and acts transitively on V. Suppose that the action is not regular, then there must be some $g \neq e$ in G that fixes some x in V. Show that g must fix all elements of V for a contradiction.

9. If $|C| \geq 3$, then we have a,b in C such that $a \neq b^{-1}$. In this case, {e, a, ab, b} is a 4-cycle in X(G, C).

24. This follows from Theorem 3.8.1. Alt(5) has 60 elements, α and β are both of order 3, and so each have 20 cycles in their action on Alt(5) by left multiplication. $\beta^{-1}\alpha = (12543)$ has order 5. Then by the theorem, X(Alt(5), {α, β}) can only be partitioned into an even number of disjoint directed cycles, and the result follows.

Chapter 4

1. Consider the vertex x = {1,…,k}. Each y with xy in E(J(2k + 1, k, 0)) is a k-subset of {k+1,…,2k+1}. So y = {k + 1,…,2k + 1}\{i}, for some $k+1 \leq i \leq 2k + 1$. Each z with yz in E(J(2k + 1, k, 0)) = {1,…,k, i}\{j}, for some $1 \leq j \leq k$.

Clearly then, 2-arcs starting at x correspond with selections of i and j, and there is a permutation that fixes {1,…,k} and {k+1,…,2k+1}, and maps any i and j to any other. Thus G_x is transitive on 2-arcs and the result follows.

3. The arcs of $X^{(s)}$ correspond with the s+1 arcs of X, and so do the vertices of $X^{(s+1)}$, so the vertex sets of $DL(X^{(s)})$ and $X^{(s+1)}$ are the same. Now observe that the edges of each graph correspond with the s+2 arcs of X.

4. If X is vertex-transitive, every connected component of X is isomorphic. So assume for contraposition that X is not arc-transitive – then the connected component with u in it will not be arc-transitive but will be vertex-transitive. By Lemma 4.3.2, it is therefore vertex-regular, so $|G_u| = 1$.

5. Since the Petersen graph is J(5,2,0), it is easy to see that any two 3-arcs can be labeled as {ab, cd, ae} and {AB, CD, AE}, and that a permutation of Ω exchanges them.

6. Since the Petersen graph is edge-transitive, we can draw the edge of interest as a spoke without losing generality. Each adjacent vertex in the outer cycle can now be included in a 1-factor in two different ways. A little drawing shows that two of the four options yield perfect matchings and two do not. Since each vertex can be used in 3 different edges to generate 2 unique 1-factors each, there are 6 1-factors (each can be obtained starting from any vertex).

Now, these 6 1-factors are equivalent under automorphism because of edge-transitivity and because the 2 1-factors generated through any edge can be seen to be equivalent. Finally, since every 1-factor is equivalent under automorphism to the spokes, no two disjoint 1-factors can be found.

7. Using the abbreviated representation from page 68, we have:

Petersen graph: {3, 2; 1, 1}.
Coxeter's graph: {3, 2, 2, 1; 1, 1, 1, 2}.
Tutte's 8-cage: {3, 2, 2, 2; 1, 1, 1, 3}

10. The Latin square graph of a group G is the Cayley graph X (G', C) where G' is the direct sum of G and H, H is the group on the same set as G with $a*_H b = b*_G a$, and C = { (h, e) for all h in H\{e}, (e, g) for all g in G\{e}, (g, g^{-1}) for all g in G }.

11. An easier way is to notice that the Latin square graph of $(Z_2)^2$ contains 4 elements at a pairwise distance of 2, while the Latin square graph of Z_4 does not.

12. For any two vertices a, b at a distance of three in Coxeter's graph, consider the distance partition from a. Any automorphism fixing a and b fixes the unique path from a to b, and hence fixes one of the neighbours of a. Now, b is connected to another vertex also at a distance of three from a, call it b'. Clearly, b' is also fixed because it is uniquely determined by a and b, and so the unique path from a to b' is also fixed. This path fixes another neighbour of a (if it were the same neighbour we would have a 5-cycle), so all of the three neighbours of a are fixed, and the automorphism must be the

identity. The conclusion follows because any two 4-arcs that agree in the initial 3-arc will not be equivalent under any automorphism.

13(*). Let G be a distance transitive graph with girth at least five. Let k = 1. Then G is at least k-arc transitive.

Consider any two k+1-arcs (they may be taken to start from the same vertex x because G is vertex-transitive). If they have an edge in common, then k-arc transitivity shows that the required automorphism exists. If not, since the end-point of each k+1 arc is at a distance of k+1 from x, there is an automorphism that exchanges them and fixes x. A little thought shows that this automorphism also exchanges the neighbours of x on each k+1 arc (using the fact that the girth is at least 2k+1), and k-arc transitivity does the rest.

For any n > 1, if the girth of G is at least 2n + 1, this argument can be applied inductively on k, to show that G is n-arc transitive (k-arc transitivity is directly implied by distance transitivity in the base case k=1, thereafter is given by the inductive hypothesis).

14. If G is an s-arc transitive graph with girth 2s-1, we have:

- There is a 2s-1 cycle
- So there is an s-arc in a 2s-1 cycle
- So there is an s-arc in a unique 2s-1 cycle (if the s-arc is in two 2s-1 cycles, then there is a 2s-2 cycle)
- By s-arc transitivity, all s-arcs lie in a unique 2s-1 cycle.
- Diameter is s-1 (imagine two vertices at a distance of s, they are joined by an s-arc, hence belong to a 2s-1 cycle, hence are joined by an s-1 arc).

Since G is s-arc transitive and has diameter s-1, it is immediate that G is distance transitive.

Chapter 5

1. The degenerate projective planes consist of a point p, through which all of the lines pass, and a line L on which all of the points lie. Every other point is trivial (lies only on L) and every other line is trivial (passes only through p).

2. These are the same as the degenerate projective planes!

3. Part I: Let S be the set of points and lines of I that are fixed by G. Then for p, q in S: $p^G = p$ and $q^G = q$. For any σ in G, we have:

(p, L) is in I iff (p^σ, L^σ) is in I iff (p, L^σ) is in I.

Therefore whenever p, q are co-linear in S, they are co-linear in I as well, and the result follows.

5(*). Let x be a vertex in a generalized hexagon of order (2, 2). Clearly x has exactly 3 neighbours. For $i < 6$, if there is an edge between two vertices at a distance i from x or two edges between a vertex at a distance i from x and vertices at a distance $i - 1$ from x, then there would be a cycle shorter than 12. So we know the intersection array is $\{3, 2, 2, 2, 2, 2; 1, 1, 1, 1, 1, y\}$.

It remains to show that $y = 3$. Each vertex u at a distance of 6 from x is connected to at least one vertex at a distance of 5 from x, call this vertex v_1. Then there are other vertices v_2 and v_3 at a distance of 5 from x along disjoint paths. By Lemma 5.6.4 each pair are on a cycle of length 12. Each pair must be connected to a mutual vertex at a distance of 6 from x. If these vertices are disjoint, we get a cycle of length 6, so they must all be u. Therefore $y = 3$ and the result is complete.

10. A Moore graph of valency 57 and diameter 2 has $(1 + 57 + 56*57) = 3250$ vertices. Use the same double counting method as is used in Section 5.9. The equations obtained are:

$$57c = \Sigma k_i \text{ and } c^2 - c = \Sigma (k_i^2 - k_i)$$

Still following 5.9, we see that $(3250 - c)\mu^2 - 104c\mu + c^2 + 56c = \Sigma (k_i - \mu)^2$.

Therefore the discriminant is non-positive, and can be calculated to be:

$104^2c^2 - 4(3250 - c)(c^2 + 56c) = 4c(c - 400)(c + 455)$.

The result follows.

11. First count 5-cycles. By Theorem 5.8.2, there are k(k-1) paths of length 2 from any given vertex x. Choosing one of these, then one of the other (k-1)

vertices at a distance of 2 from x gives a 5-cycle. Each 5-cycle can be generated thus in 2 ways, so there are $k(k-1)^2/2$ 5-cycles.

Now, count the number of vertex/5-cycle pairs in two ways. If n is the number of 5 cycles we have:

$(k^2 + 1)$ vertices * $k(k-1)^2/2$ 5-cycles through each = n 5-cycles * 5 vertices in each.

So $k(k-1)^2(k^2 + 1)/2$ divides 5, and the result follows easily.

Chapter 12

1. That $\rho_a\rho_b = \rho_b\rho_a$ if $\langle a\rangle = \langle b\rangle$ or $a\cdot b = 0$ is easily seen by drawing a picture. Conversely, expanding $\rho_a\rho_b(x)$ and $\rho_b\rho_a(x)$ using the formula $\rho_h(x) = x - 2(x\cdot h/h\cdot h)h$ and simplifying yields the equation $(a\cdot b)(x\cdot b)a = (a\cdot b)(x\cdot a)b$. This forces either $a\cdot b = 0$ or $(x\cdot a)b = (x\cdot b)a$. Pick any x orthogonal to a, and it must also be orthogonal to b, and vice versa – this forces $\langle a\rangle = \langle b\rangle$.

2. That D_n is indecomposable is given as Lemma 12.4.1. Any two non-orthogonal vectors in D_n will have inner product ±1. Since we are interested only in the lines spanned by the vectors we can assume that $a\cdot b = -1$. Then the third line in the star can be seen to be $- a - b$, and it is trivial to check that this is also in D_n.

3 & 4(*). One way to do this is by working out the distance partition from a vertex x (note the graph may not be distance regular, so this is not identical for every vertex). Wolog take x +ve.

Note that ±y is a neighbour of x if the angle between $\langle x\rangle$ and $\langle y\rangle$ is $\pi/3$. Again wolog take y –ve, so that y is a neighbour of x. Now, it can be seen that $- x + y$ belongs to L (star-closure), and is a neighbour of y and –x. So far, then, we have {x, N(x), N(-x), -x}, and the one to one correspondence between N(x) and N(-x) required for exercise 4 is apparent. Vectors orthogonal to x clearly must be orthogonal to –x, but orthogonal to some member of N(x) or N(-x) because L is decomposable. So they are each at a distance of 2 or 3 from x. Thus the diameter of Y is 3 as required.

5. Since the members of S are in A_n, we will certainly have 2's on the diagonal of the Gram matrix, and non-negative inner products guarantees it

will be $2I + A(X)$ for some X. If it satisfies $B^TB = 2I + A(X)$ where B is the incidence matrix of some graph Y we will have that $X = L(Y)$. First, note that since all inner products are non-negative, replacing all -1s with 1s in the vectors of S will have no effect on the Gram matrix. Now we need only deduce the characteristics of Y. This is done by noticing that if an odd number of distinct elements of A_n have a position in common (which corresponds to their representing adjacent edges in Y) then the inner product of one of the pairs has to be odd. Thus Y has no odd cycles, and is bipartite.

7(*). Let X be a graph with minimum eigenvalue at least -2. From Theorem 12.8.1, we need only show that $A(X) + 2I$ is not the Gram matrix of a set of vectors contained in E_8. Imagine that there is an independent set of size ≥ 9 in X. Wolog index 9 of the independent vertices $\{1,\ldots,9\}$. Then the top left corner of $A(X) + 2I$ will be $2I_9$ and this implies that if the set of vectors with Gram matrix $A(X) + 2I$ is labelled $\{v_1,\ldots\}$ then $v_i \cdot v_j = 0$ for $i \neq j \leq 9$. It is easy to see that no such 9 vectors exist in E_8, and so X is a generalized line graph.

Chapter 13

1. Proceed by induction on the size of the square submatrix S. The base case is that S is 2*2. In this case there is a 0 in S, because the two columns of D passing through S cannot have 1's in the same rows. The inductive step has two cases: if there is a column in S with zero or one non-zero entries, then a cofactor expansion along that column gives the result. Otherwise, $S^T\mathbf{1} = 0$, so 0 is an eigenvalue of S^T, S is thus non-invertible and det S = 0 as required.

4. The eigenvalues of K_m and K_n are m^{m-1}, 0 and n^{n-1}, 0. By indexing the vertices of $K_m + K_n$, we can see that $Q(K_m + K_n)$ can be split into quadrants that are $Q(K_m)$, 0, 0 and $Q(K_n)$. Then any eigenvector of $Q(K_m + K_n)$ has its first m positions an eigenvector of K_m and its last n an eigenvector of K_n. Thus every eigenvalue of K_m or K_n is also and eigenvalue of $K_m + K_n$. Now, applying Lemma 13.1.3, we can compute the eigenvalues of $K_{m,n}$, and thus its characteristic polynomial.

6. By Lemma 13.1.5, we can write the given expression as $nx^TQx/x^TQ(K_n)x$. Then $Q(K_n) = Q(X) + Q(K_n \backslash X)$, so we can rewrite the denominator. Then using Lemma 13.1.3 and putting x a non-constant eigenvector shows that this expression is equal to the corresponding eigenvalue. Finally, it remains to see that the minimum over such x is the minimum over all x.

7. The first part follows from Lemma 13.5.2. Equality if T is a star follows from Exercise 4. Inequality if T is not a star follows from the proof of Theorem 13.5.1, by taking S to be a single vertex (when T is not a star there will always be some edges with no ends in S).

Chapter 15

1. Imagine that A + {x} is dependant for every x in B\A. Then the submodular property gives rk(A + x + y) + rk(A) ≤ rk(A + x) + rk(A + y) = 2rk(A). So rk(A + x + y) = rk(A) and clearly an induction leads us to conclude that rk(A + B) = rk(A) which is a contradiction because B is a subset of A + B and has rank greater than A.

2. Apply the submodular property to C and D. Then since rk(C∩D) = |C∩D|, rk(C) ≤ |C| - 1, rk(D) ≤ |D| - 1, and |C+D| = |C| + |D| - |C∩D|, we have rk(C + D) ≤ |C+D| - 2. This implies that (C+D)\{x} is not independent, and thus that it contains a circuit.

3(*). It is immediate that inclusion is partial order on the set of flats. An intersection of flats is a flat, so for any set S of flats in a matroid M the intersection of S is a greatest lower bound. Define f(A) to be the subset of Ω containing all members of A, and also all elements in circuits C such that A contains |C| - 1 elements in C. Then f(union of all members of S) = U is a lowest upper bound for a set S of flats. It is clear that U is an upper bound. Now, for any subset V of U, consider U\V.

Case 1: U\V includes some x not in a circuit in U, and hence in s for some s in S, so that s is not contained in V.

Case 2: U\V includes some x that belongs to a circuit in U, in which case it must contain another y from the same circuit (or V wouldn't be a flat). But then either x or y must be in some s in S, with s not contained in V.

5(*). Let A be the complement of the support of a word w in the rowspace of D. rk(A) is the rank of the matrix **A** with rows that are the restriction of the rows of D to A. rk(A + x) for any x not in A is the rank of the matrix **A + x** with rows that are the restriction of the rows of D to A + x. Since a linear combination of the rows of D gives w, and w is 0 on A and non-zero on x, a linear combination of the rows of **A + x** gives (0,…0,k) for some k, and we can row reduce **A + x** to a matrix with A in its top left hand corner, and a 1 in its bottom right hand corner. Clearly then rk(A + x) = rk(A) + 1, as required.

8. Firstly, if e is a loop, there are no acyclic orientations.

If e is a cut-edge, then assume that e = st, and that v is in the same connected component of X\e as s. Then the other component of X\e containing t must be acyclic, and so must have a source. This source must therefore be t, and each orientation of X/e with v as unique source can be turned into an orientation of X with v as the unique source in exactly one way by orienting e s -> t.

If e is neither a loop nor a cut-edge, there are two more cases. If e = st and neither s nor t is v, then the result follows exactly as in Theorem 15.6.1. If wolog s = v then still proceed by counting acyclic orientations of X\e with v as the unique source. Now, there are κ(X/e,v) of these with no directed path from s=v to t, but when we re-add e, we have all edges leaving t except for the one from v. Reversing each of the edges leaving t gives another acyclic orientation. So we have 2κ(X/e,v) orientations. Finally, there are still κ(X\e,v) - κ(X/e,v) more orientations to count (those derived from acyclic orientations of X\e with a directed path from v to t).

κ(X,v) = R(M(X);0,-1), and there are 714 such orientations in the Petersen graph.

9. Let e = uv. Then consider any homomorphism f:X -> Y. There is a 1-1 correspondence between homomorphisms where f(u)f(v) is an edge in Y and homomorphisms from X/e to Y. There is also a 1-1 correspondence between homomorphisms where f(u)f(v) is not an edge in Y and homomorphisms from X\e to Y. The result follows.

11. Noting that rk(M\e) = rk(M) whenever e is not a coloop, that rk(M/e) = rk(M) - 1 whenever e is not a loop, and that the same element is never a loop and a coloop, the result is an easy induction on |M|.

12(*). First, notice that any non-zero value can be assigned to a loop and that no non-zero value can be assigned to a bridge, so that the first two parts of the equation are easy.

Let e = uv. Clearly every nowhere-zero q-flow on X can be used to generate one on X/e (simply assign the same values to all the edges). All nowhere-zero q-flows on X/e can be generated in this way except those where the total on edges between uv and N(u) = the total on edges between uv and N(v) = 0 (these could only be generated by flows on X where the value on e was 0). However, each of these flows corresponds with a nowhere-zero q-flow on X\e, and the result follows.

13. $F(X, q) = (-1)^{|E(X)| - |V(X)| + c} R(M(X), -1, -q)$