Summer Research 2013
James Madison University Deelan Jalil

Integer Invariants of Abelian Cayley Graphs

Introduction.

This paper is an informal summary of the trials and tribulations experienced by Dr. Joshua Ducey and I as we researched ways to predict the Smith normal form (SNF) of various incidence matrices. We'll define and discuss multiple terms and concepts throughout this paper.

The title of our project is *Integer Invariants of Abelian Cayley Graphs*. Let me explain what the title means. First is the 'integer' part. This simply means that every number we worked with was an integer. We limited ourselves and were not allowed to use fractions or irrational numbers. Second, an 'invariant' is a property that doesn't change; it doesn't vary. Smith normal form, is an invariant. That is, regardless of the way we order the rows and columns of our matrix, once we perform elementary row and column operations to get it into Smith normal form, we always get the same SNF. Lastly, the graphs we dealt with were all abelian Cayley graphs.

Define a graph to simply be a collection of vertices that are connected by edges. Given this very lenient definition, the graphs can get very messy and erratic, as you can imagine. However, when we limit ourselves to working with abelian Cayley graphs, we see that they are very pretty and highly symmetric. So what is an abelian Cayley graph, you might ask. Good question. Abelian Cayley graphs consist of two parts, an abelian group and a connecting set. An abelian group G is defined to be a set of elements under some operation that is commutative. Furthermore, in order for it to be a group, the conditions of closure, associativity, existence of an identity, and existence of inverses must be satisfied. The second component of an abelian Cayley graph is the existence of a connecting set. A connecting set E is any subset of the group. You use this set E to determine adjacency. Let $g, h \in G$. Then g and h are adjacent if $gh^{-1} \in E$. The first type of abelian Cayley graph we studied comes from the Hamming association scheme.

Hamming Distance and Hamming Association Schemes.

Let X be the set of all vectors of length n, with coordinates taken from some alphabet of size q. Define the *Hamming distance* $d(x, y)$ between $x, y \in X$ to be the number of coordinates in which x and y differ. We say $''x$ *is* k*-incident to* y'' if and only if $d(x, y) = k$, for $0 \leq k \leq n$. For example, the Hamming distance between APPLE and SMILE is 3 since the two words differ in the 3 positions. The Hamming distance between between MATH and COOL is 4. The distance between $(0, 1, 0, 1, 1)$ and $(1, 1, 0, 0, 1)$ is 2.

The set X of all such vectors with length n from an alphabet of size q as you run through the all the possible distances k where $0 \leq k \leq n$ is called the Hamming association scheme.

Association Schemes.

An association scheme consists of a set X together with a partition R of $X \times X$ into $n + 1$ binary relations $R = \{R_0, R_1, R_2, ...R_n\}$ which satisfy:

1. $R_0 = \{(x, x) : x \in X\}$ which is called the identity relation.
(When $k = 0$, everyone is related to himself.)

2. For $i = 0, 1, 2, ..., n$, the inverse $R_i^{-1} = \{(y, x)|(x, y) \in R_i$ of the relation R_i belongs to S.
(The matrix is symmetric; if x is distance k from y, then y is distance k from x.

3. If $(x, y) \in R_k$, the number of $z \in X$ such that $(x, z) \in R_i$ and $(y, z) \in R_j$ is a constant C_{ijk}. (You are counting the number of $z \in X$ that are distance i away from x and distance j away from y, while $d(x, y) = k$.)

Quick Sidenote:
To calculate the constant that goes in front of each R_k, you have to consider each of the possible $d(x, y) = k$, where $0 \leq k \leq n$. You do each distance k case separately. First, construct x and y for each individual cases of k. Next, count how many z's you can construct such that z is distant i away from x and distance j away from y. It helps to keep tallies of the differences when trying to count how many possible such z's you can make. Note that x will have i tallies and y will have j tallies.

Back to the Hamming Association Scheme.

To see that this is an abelian Cayley graph, simply observe that we have our abelian group $X = Z_q \times Z_q \times \cdots \times Z_q$ (n times) and our connecting set is now E_k, which we'll define to be the set of tuples that have exactly k components that are not the identity.

We can pick a distance k and encode this association (the set X and the chosen incidence relation) into a matrix. Note that the size of X is q^n because you have q choices for each of the n positions. When we create our matrix, we index the rows and columns by the set X. If the relation k is satisfied between two vectors, we denote it with the entry 1; otherwise, we use 0. We repeat this process over and over again until we completely fill out our matrix. That matrix is then called an adjacency or an incidence matrix.

Below is a pictorial representation for the adjacency matrix formed for distance $k = 1, 2, 3$ for the set of length 4 vectors constructed from an alphabet of size 3. White pixels denotes a zero-entry, that the two elements are not incident; black pixels denote a value of one, that the two elements are indeed incident.

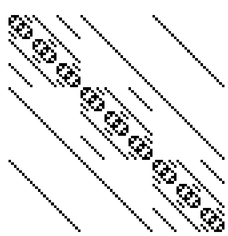

81×81 adjacency matrix for $k = 1$ in $H(4, 3)$

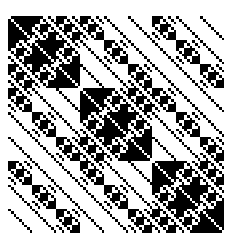

81×81 adjacency matrix for $k = 2$ in $H(4, 3)$

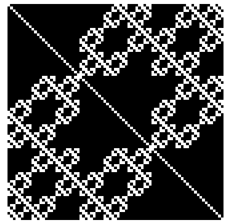

81×81 adjacency matrix for $k = 3$ in $H(4, 3)$

No matter how the adjacency matrix A is formed, the eigenvalues, up until sign, are always the same. So what is an eigenvalue? If you think all the way back to your linear algebra course, somewhere along the way you learned that a number is an eigenvalue if and only if there exists a nonzero vector v such that $Av = \lambda v$. Then λ is an eigenvalue and v is an eigenvector.

Finding the eigenvalues of a matrix.

By definition, an eigenvalue λ satisfies $Av = \lambda v$, where v is a nonzero vector. This can be rewritten as $Av = \lambda I v$, where I is the identity. Now observe that $0 = \lambda I v - Av = (\lambda I - A)v$. So $det(\lambda I - A)$ must equal zero to satisfy the equation, since v is a nonzero vector. Now solve for λ by creating the matrix $\lambda I - A$, finding its determinant (which will have λ's in it), and setting it equal to 0. This will give you your eigenvalues λ.

So we view the eigenvalues as expressing (or being) some fundamental invariant of the incidence structures. Invariant means it doesn't change. So for all these incidence matrices that we're talking about (that come from fixing the set X and the particular distance to define the relation), their eigenvalues are known. There is a formula for it. But what is not known is the Smith normal form of these matrices. Our goal is to figure out what the Smith normal form of every matrix is depending on how you fix your parameters n , q, and k.

Smith Normal Form.

Ok, so what is Smith normal form? We keep talking about it but have yet to define it. Given any integer matrix A, you can find square integer matrices P, Q with determinant ± 1 such that:

$PAQ = \begin{pmatrix} d_1 & 0 & 0 & \cdots & 0 \\ 0 & d_2 & 0 & \cdots & 0 \\ 0 & 0 & d_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & d_n \end{pmatrix}$ where $d_1 | d_2$, $d_2 | d_3$ ... $d_{n-1} | d_n$. This diagonal matrix is called the Smith normal form of A. Note that $d_1, d_2, ..., d_n$ are called invariant factors. Furthermore, if you break up the invariant factors into primes and their powers, those primes and their powers are called the elementary divisors of the SNF.

Every integer matrix has a Smith normal form. One application of SNF is that it can help us distinguish non-isomorphic graphs. Take the set of all $n \times n$ matrices. In here we'll have adjacency matrices, regular integer matrices, and matrices that are in Smith normal form. When we partition them up into equivalence classes, each class will only have one matrix in SNF. And here's the good stuff. If you take two graphs that are isomorphic and construct their adjacency matrices, it turns out that they'll fall into the same equivalence class because they'll have the same SNF! How cool is that! But note that this does not work the other way around. Thus, if two matrices have the same SNF, it does not mean you can conclude that their graphs are isomorphic.

Comparing the product of eigenvalues to the product of invariant factors

Under certain conditions (when the matrix is symmetric), we can diagonalize A. The diagonal entries will be the eigenvalues of the matrix. Because we can diagonalize it, it means there exists some invertible matrix P (can have any determinant that is not zero) such that

$P^{-1}AP = \begin{pmatrix} \lambda_1 & 0 & 0 & \cdots & 0 \\ 0 & \lambda_2 & 0 & \cdots & 0 \\ 0 & 0 & \lambda_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda_n \end{pmatrix}$. Note that the $det(P^{-1}AP) = det(P^{-1})det(A)det(P)$.

Furthermore, since A has integer entries, we can put it in SNF by finding some R and Q with determinant ± 1 such that

$RAQ = \begin{pmatrix} s_1 & 0 & 0 & \cdots & 0 \\ 0 & s_2 & 0 & \cdots & 0 \\ 0 & 0 & s_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & s_n \end{pmatrix}$. For this matrix, $det(RAQ) = (s_1)(s_2)(s_3)...(s_n)$.

Thus we can conclude that $det(P^{-1}AP) = det(RAQ)$, up until sign. They are the same up until sign because the determinant of P^{-1} and P cancel out leaving A, and R and Q have determinant ± 1, leaving A as well.

This is important because it is telling us that from the determinant, which is simply the product of the eigenvalues, we can pull out primes and prime powers. From those primes and prime powers, we can identify which primes are involved in the SNF matrix.

Making the program and analyzing the data.

So we wanted to come up with a technique or formula to compute the Smith normal form of any matrices regardless of the chosen parameters. To do so, Dr. Ducey wrote a program in Sage that allowed you to input n, q, and k. It would create the vectors based on n and q and then put the appropriate entry (zero or one) in the matrix based on your Hamming distance k.

The program then converted the matrix into Smith normal form and it would output the invariant factors and their multiplicities. We collected our data by fixing different n's and q's and running through all of the possible k's. We tried looking for some underlying pattern but that was difficult and we were unsuccessful.

We then decided to compare the invariant factors (and their multiplicities) against the eigenvalues (and their multiplicities). In comparing them, we saw some correlation, but there were gaps. Because of those gaps, we decided to break up both the invariant factors and the eigenvalues into primes and power of primes and compare those multiplicities instead. But again we ran into problems because not everything was lining up.

As we further analyzed the data, we saw that if a prime p did not divide q, then that prime and its powers would occur the same number of times in both the eigenvalues and as elementary divisors. Finally, we were making progress!! We could partially determine the elementary divisors of the SNF! However, when p divides q the multiplicities didn't match up. It turns out that there is very little known about how to predict any information about the occurrence of primes as an elementary divisor when $p \mid q$.

When $p \nmid q$, we can use eigenvalues to predict the occurrence of p as an elementary divisor.

Eigenvalues and their multiplicities are known for each matrix A_k that we construct from our parameters n, q, and k; someone has come up with a formula for these. The determinant of each A_k is also known. The elementary divisors come from the primes that divide $det(A_k)$. Let these primes be p_i^j, where p_i is a prime $2, 3, 5, 7, etc...$ and j tells you what power that prime is, so $2^3, 3^8, 5^2, 7^3, etc....$ Let $e_{p_i^j}$ denote the multiplicity of each prime power. We can prove a formula for $e_{p_i^j}$ when $p_i \nmid q$.

In other words, we are stating that if a prime shows up in the eigenvalues that does not divide q, then it's occurrence will be the same in the Smith normal form as it was in the eigenvalues. Our argument involves conjugating A_k with the matrix M, which will produce a diagonal matrix.

There exists a character table matrix M over the ring of p-adic integers $\mathbb{Z}_p$, such that $MA_kM^{-1} = \begin{pmatrix} y_1 & 0 & 0 & \cdots & 0 \\ 0 & y_2 & 0 & \cdots & 0 \\ 0 & 0 & y_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & y_n \end{pmatrix}$, where y_a are the eigenvalues of A_k. The operation MA_kM^{-1} diagonalizes any linear combinations of A_k. When we conjugate any matrix with its character table, the entries produced along the diagonal are its eigenvalues. In fact, this is the same diagonal form in the sense of Smith normal form since M is unimodular over $\mathbb{Z}_p$. Thus, the multiplicity of p^i as an elementary divisor is the number of entries along the diagonal y_a exactly divisible by p^i.

Character tables.

Let $G = (C_q)^n$, where $C_q =< x >$ is the cyclic group of order q. Let $E_k := \{g \in G|$ g has exactly k components not equal to the identity$\}$. Let M be the character table of G with columns ordered in the same way as for A_k.

Quick Sidenote:
A character of an abelian group G is a homomorphism $\chi : G \to \mathbb{C}$. Note that the trivial character of G is the homomorphism $\mathbf{1}_G$ defined as $\mathbf{1}_G(g) = 1$ for all $g \in G$. It turns out that the number of characters is $|G|$, so M is a square matrix. Furthermore, by known orthogonality relations, the dot product of two rows will always be zero and the dot product of a row with itself will always be the order of the group.

Define $\overline{M}^t$ to be the complex conjugate transpose matrix of M. In other words, you first take the complex conjugate of each element; that is, $a+bi$ becomes $a-bi$. Then you transpose the matrix, so the rows of M become the columns of $\overline{M}^t$. Notice that $M\overline{M}^t = |G|I$; in other words, the product of M and $\overline{M}^t$ is a diagonal matrix with the order of G as its diagonal entries. It follows, $M\frac{1}{|G|}\overline{M}^t = I$.

Let $M^{-1} = \frac{1}{|G|}\overline{M}^t$. Then $MA_kM^{-1} = \begin{pmatrix} \sum\limits_{x\in E} \chi_1(x) & 0 & \cdots & 0 \\ 0 & \sum\limits_{x\in E} \chi_2(x) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sum\limits_{x\in E} \chi_{|G|}(x) \end{pmatrix}$

If M and M^{-1} were unimodular, we'd be done; the diagonal entires would contain the elementary divisors of the SNF. Thus, to satisfy the unimodular conditions, we'll work over $\mathbb{Z}_p$, where $p \nmid |G|$. When we do so, the only numbers that are not invertible are powers of p. Therefore, we know M and M^{-1} are unimodular. So by working over $\mathbb{Z}_p$, we were able to conclude M and M^{-1} were unimodular. Thus for primes $p \nmid |G|$, we can get the elementary divisors.

Proving that MA_kM^{-1} diagonalize A_k, as claimed above.

We want to prove that MA_kM^{-1} does in fact diagonalize any matrix A_k. Let $G = (C_q)^n$, where $C_q =< x >$ is the cyclic group of order q. Let $E_k := \{g \in G|$ g has exactly k components not equal to the identity$\}$. Now define $g \sim h$ if and only if $gh^{-1} \in E_k$. Note that A_k is the adjacency matrix for the Cayley graph of G with the connecting set E_k. Let M be the character table for G, with columns in the same order as those of A_k. We claim:

$$MA_kM^{-1} = diag(\sum_{y \in E} \chi(y)). \tag{1}$$

We'll first compute MA_k. Consider the (χ, g) entry of MA_k. Observe that it is $\sum\limits_{h \in G | h \sim g} \chi(h)$. This is because A_k is made up of all zeros and ones. So when hitting a row of M by a column of A_k the only elements that end up surviving are the row elements $\chi(h)$ that correspond to a 1 in the column of the adjacency matrix. This happens when $h \sim g$. Summing over them gives us the equation described.

Now consider the (χ, ψ) entry of MA_kM^{-1}. It is $\sum\limits_{g \in G} \sum\limits_{g \in G | h \sim g} \chi(h)\overline{\psi(g)}$. This is because you are summing over the product all of the g entries in the rows of the first matrix by $\psi(g)$ entries down the columns of the second matrix. The bar sign signifies that they are column entires.

Finally, we'll see that this double sum produces a diagonal matrix. Observe the following set of equalities:

$$\begin{aligned}
\sum_{g \in G} \sum_{g \in G | h \sim g} \chi(h)\overline{\psi(g)} &= \sum_{g \in G} \sum_{e \in E_k} \chi(eg)\overline{\psi(g)} \\
&= \sum_{g \in G} \sum_{e \in E_k} \chi(e)\chi(g)\overline{\psi(g)} \\
&= \sum_{g \in G} \chi(g)\overline{\psi(g)} \sum_{e \in E_k} \chi(e) \\
&= \sum_{e \in E_k} \chi(e) \sum_{g \in G} \chi(g)\overline{\psi(g)} \\
&= \begin{cases} \sum\limits_{e \in E_k} \chi(e) & \text{if } \chi = \psi \\ 0 & \text{if } \chi \neq \psi \end{cases}
\end{aligned} \tag{2}$$

Linear transformations and p-adic integers.

The matrix A_k can also be viewed as defining a linear transformation $f : \mathbb{R}^n \to \mathbb{R}^n$. If we know where the basis of the domain gets sent to, the whole linear transformation is defined. Once you get the basis of the domain and send it to the basis of the range, it's a linear transformation. We can do this under the p-adic integers $\mathbb{Z}_p$, which contains and is larger than $\mathbb{Z}$, but only has one prime, namely p.

Suppose an adjacency matrix had the following SNF: $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 6 & 0 \\ 0 & 0 & 0 & 60 \end{pmatrix}$

Note that this can also be rewritten as: $\begin{pmatrix} 1*1*1 & 0 & 0 & 0 \\ 0 & 1*3*1 & 0 & 0 \\ 0 & 0 & 1*3*2 & 0 \\ 0 & 0 & 0 & 5*3*2^2 \end{pmatrix}$

Thus, we would view the matrix as $1*1*2*2^2$ in $\mathbb{Z}_2$, $1*3*3*3$ in $\mathbb{Z}_3$, and $1*1*1*5$ in $\mathbb{Z}_5$.

Zooming in further.

It seemed that the scope of what we were doing was too large and preventing us from making much progress. We therefore decided to only concentrate on one specific case, the case when $q = 2$, so $H(n, 2)$. If we could predict the SNF of those matrices, maybe we could extent the argument to different parameters. In doing so, we slowly began noticing some patterns and underlying structures with fixing q and running through the different cases of n and k.

First, we observed that A_k and A_{n-k} had the same exact SNF! Furthermore, they have the same eigenvalues up until sign. This is proved below! Thus, if we could predict the SNF for the first half of the data, we would know what it is for the remaining half when q is fixed to be 2.

When we fix both $q = 2$ and $k = 1$, the following happens:

1. The eigenvalues for $(n, 2, 1)$ ranges from $-n$ to n, increasing by 2 each time. So the eigenvalues are $-n + 2l$ for $0 \leq l \leq n$.
2. Furthermore, the multiplicity of the eigenvalue is $\binom{n}{l}$.
3. When n is odd, we can predict the elementary divisors and their multiplicities from the eigenvalues since 2, which divides the order of the group, will never appear as an elementary divisor.
4. When n is even, the multiplicity of 2^i as an elementary divisor is equal to the number of eigenvalues exactly divisible by 2^{i+1}.
5. The highest invariant factor in the Smith normal form value with $q = 2$ and $k = 1$ will have multiplicity 2.

Note that items 4 and 5 above are conjectures.

As we further analyzed those cases, we stumbled on the fact that when we fix $k = 1$, we are looking at the n-cube adjacency graph, which consists of all tuples of 0 and 1 which are connected when their distance is one. Therefore, the case when $n = 3$, $q = 2$, and $k = 1$ is essentially trying to figure out the Smith normal form of a cube at the origin of the 3-D graph.

Proving the matrices A_k and A_{n-k} for the n-cube have the same Smith normal form.

Fix $q = 2$ and consider the matrix A_k. This is a $q^n \times q^n$ matrix with the rows and columns indexed by the set X, which consists of all possible length n vectors with entries 0 or 1. We define $x, y \in X$ to be associates if the Hamming distance between them is k. We denote this relation by the entry 1 in the $x - y$th position. Next, consider the permutation mapping R defined by $R : y_b \to \vec{j} - y_b$. (Note that $\vec{j}$ is the all 1 vector.) This permutation sends each row to its complement.

Let P be the permutation matrix for the function R. Note that P is also $q^n \times q^n$ and contains exactly a single 1 entry along each row and column, since it permutes rows of I based on the permutation defined above. Note that $det(P) = \pm 1$. We will see that $PA_k = A_{n-k}$.

We start with our matrix A_k. Every column indexed by x_a will have z 1 entries corresponding to y_b, denoting their Hamming distance to be k for $1 \le a, b \le q^n$. Once the permutation P is applied to A_k, the entires along row y_b will now correspond to row $\vec{j} - y_b$.

Therefore, since the distance between x_a and y_b was k, the distance between x_a and $\vec{j} - y_b$ is now $n - k$, and is marked with the entry 1. This process builds up the matrix A_{n-k}, so $PA_k = A_{n-k}$. And since P is unimodular, A_k and A_{n-k} will have the same Smith normal form and the same eigenvalues up till sign.

<u>We can get the SNF for odd n-cubes from the eigenvalues.</u>

When $q = 2$, elementary divisors and their multiplicities are known for $p \nmid q$. Thus the only case we need to consider is when $p \mid q$; that is, when $p = 2$. The statement below regarding eigenvalues shows that when n is odd, no powers of 2 appear in the eigenvalue or the elementary divisors. Therefore, the only case left to consider after this proof is when n is even.

Let A_n denote the adjacency matrix for the n-cube graph Q_n. Note that A_n is the distance one association matrix inside $H(n, 2)$. First, consider the case when n is odd. Then, A_n has eigenvalues $-n + 2l$ for $0 \le l \le n$, with multiplicity $\binom{n}{l}$. Clearly in the odd case, zero is not an eigenvalue and all of the eigenvalues are odd. In other words, none of the eigenvalues are divisible by 2. The product of the eigenvalues is equal to the product of the elementary divisors of A_n. Therefore, since the product of the eigenvalues are odd, we can conclude none of the elementary divisors are divisible by 2. Hence, the SNF can be fully predicted from the eigenvalues.

<u>Rough outline for official paper.</u>

At about this point, we started thinking about how to organize our official fancy math paper. Our initial ideas were the following:

Abstract: Include a 3-5 sentence summary of what our results were and what they can be applied to.

Intro: Define terms and give information about what others have done in the field so far.
Main Result: Conjugating any adjacency matrix A_k by the group's character table will diagonalize it and yield the information necessary to construct the SNF of that adjacency matrix.

Applications:

1. *Hamming scheme* $H(n, q)$

Let our abelian group be $G = \mathbb{Z}_q \times \mathbb{Z}_q \cdots \times \mathbb{Z}_q$ (n times) and define the connecting set E_k be defined to be $E_k = \{g \in G |$ g has exactly k components not equal to the identity$\}$. The adjacency matrix of these graphs as $0 \le k \le n$ are the association matrices of the Hamming scheme. We have results pertaining to the SNF of these matrices. Note that our results apply to any integral matrix in the Bose-Mesner Algebra of $H(n, q)$.

2. *The n-cube and Cartesian products of complete graphs*

Let $G = \mathbb{Z}_2 \times \mathbb{Z}_2 \times \cdots \times \mathbb{Z}_2$ n times and set $E_k = E_1$. This graph is the n-cube. When n is odd, the case is solved; we know what the SNF will be. However when n is even, we run into trouble when predicting the 2-part. But we do have a conjecture that we strongly, strongly believe is true. Furthermore, we can recover results of H. Bai.

More generally, for the Cartesian product of complete graphs, let $G = \mathbb{Z}_{q_1}^{c_1} \times \mathbb{Z}_{q_2}^{c_2} \times \cdots \times \mathbb{Z}_{q_n}^{c_n}$ and $E_k = E_1$. We can recover the results of Jacobson, Niedermair, and Reiner. We get the same formula they do but in a much simpler way.

3. *Laplacian matrix L*

The Laplacian matrix is defined to be $L = dI - A$, where d is the degree of each vertex, the number of edges coming out from the vertex, and A the adjacency matrix A_1. We diagonalize the Laplacian matrix by conjugating it with it's character table.

$$\begin{aligned} L &= dI - A \\ M(L)\frac{1}{|G|}\overline{M}^t &= M(dI - A)\frac{1}{|G|}\overline{M}^t \\ &= MdI\frac{1}{|G|}\overline{M}^t - MA\frac{1}{|G|}\overline{M}^t \\ &= dI - MA\frac{1}{|G|}\overline{M}^t \\ &= dI - diag(|G| \sum_{e \in E_k} \chi(e)) \end{aligned} \tag{3}$$

Notice that M also diagonalizes $dI + cA$. Thus when the Cayley graph is regular, say when the connecting set satisfies $E = E^{-1}$, so our results apply to the Laplacian, signless Laplacian, etc...

Calculating eigenvalues $\sum_{e \in E_k} \chi(e)$.

Let $G = \mathbb{Z}_q \times \mathbb{Z}_q \times \cdots \times \mathbb{Z}_q$ (n times) and define $E_k = \{g \in G|$ g has exactly k components not equal to the identity$\}$. Earlier we saw that $MA_k \frac{1}{|G|}\overline{M}^t = diag(\sum_{e \in E_k} \chi(e))$, where χ is an irreducible character of G. Note that χ is of the form $\chi(a_1, a_2, ..., a_n) = \chi_1(a_1)\chi_2(a_2)\dots\chi_n(a_n)$ where χ_i are irreducible characters of $\mathbb{Z}_q$. So we have $\sum_{e=(e_1,e_2,...,e_n) \in E_k} \chi(e) = \sum \chi_1(e_1)\chi_2(e_2)...\chi_n(e_n)$. Let $K \in \binom{[n]}{k}$, where $[n]$ is the set of numbers from 1 to n and k are the number of fixed principal character positions.

The eigenvalues $\sum_{e \in E_k} \chi(e)$ are $\sum \prod_{[i \in K]} \sum_{e_i \in \mathbb{Z}_q | e_i \neq 0} \chi_i(e_i)$.

Observe that

$$\sum_{e \in \mathbb{Z}_q | e_i \neq 0} \chi_i(e_i) = \begin{cases} q-1 & \text{if } \chi \text{ is a principal character} \\ -1 & \text{otherwise.} \end{cases} \tag{4}$$

This is because $0 = <\chi_i, 1_G> = \sum_{e \in \mathbb{Z}_q} \chi_i(e) = 1 + \sum_{e \in \mathbb{Z}_q | e \neq 0} \chi_i(e)$. This leads to the following formula for eigenvalues: $\sum_{j=0}^{l} \binom{l}{j}\binom{n-l}{k-j}(q-1)^j(-1)^{k-j}$ with multiplicity $\binom{n}{l}(q-1)^{n-l}$.

Applying this formula for the n-cube, we set $q = 2$ and $k = 1$. The sum simplifies down to $-n + 2l$ for $0 \leq l \leq n$. This is because when we set $k = 1$, the first two terms are the only ones that don't go to zero. The others go to zero because we get the invalid expression to choose a negative number, of which there are zero ways to do. Thus when only taking into account $j = 0$ and $j = 1$, the first term simplifies down to $(1)(n-l)(1)(-1)$ and the second ones simplifies down to $(l)(1)(1)(1)$. Summing over these two, we get $(-n+l)+l = -n+2l$, as desired. Also, when we set $q = 2$ into the multiplicity formula, it clearly yields $\binom{n}{l}$.

Modifying previous argument for more general cases.

Let $G = \mathbb{Z}_{q_1} \times \mathbb{Z}_{q_2} \times \cdots \times \mathbb{Z}_{q_n}$ and define $E_1 = \{g \in G|$ g has exactly one component not equal to the identity$\}$. Earlier we saw that $MA_k \frac{1}{|G|}\overline{M}^t = diag(\sum_{e \in E_1} \chi(e))$, where χ is an irreducible character of G. Note that, χ is of the form $\chi(a_1, a_2, ..., a_n) = \chi_1(a_1)\chi_2(a_2)\dots\chi_n(a_n)$ where χ_i are irreducible characters of $\mathbb{Z}_{qi}$. Let $K \in \binom{[n]}{k}$, where $[n]$ is the set of numbers from 1 to n and k are the fixed positions. We claim that our eigenvalues $\sum_{e \in E_k} \chi(e)$ are $\sum \prod_{[i \in K]} \sum_{e_i \in \mathbb{Z}_{qi} | e_i \neq 0} \chi_i(e_i)$.

Observe that

$$\sum_{e \in \mathbb{Z}_{qi} | e_i \neq 0} \chi_i(e_i) = \begin{cases} q_i - 1 & \text{if } \chi_i = 1_G \\ -1 & \text{if } \chi_i \neq 1_G \end{cases} \tag{5}$$

For $\sum_{e \in E_k} \chi(e)$, say $\chi_{i1}, \chi_{i2}, ..., \chi_{il}$ are principal. Where there are no principal characters, you sum up all of the -1. If a χ_{ij} is principal, you add $q_i - 1$. This gives us the following expression for eigenvalues: $\sum_{e \in E_k} \chi(e) = -n + \sum_{j=1}^{l} q_{ij}$ with multiplicity at least $\prod_{i \notin i_1, i_2, ..., i_l} (q_i - 1)$. Note that G is a regular graph with valency $\sum_{i=1}^{n} (q_i - 1) = d$. So we can use this formula to find the eigenvalues for the Laplacian Matrix.

$$n - \sum_{j=1}^{l} q_{ij} + \sum_{i=1}^{n} (q_i - 1) = \sum_{i=1}^{n} q_i - \sum_{j=1}^{l} q_{ij} = \sum_{\chi_i = 1_{\mathbb{Z}_{qi}}} q_i \tag{6}$$

In English, we pick our χ, which is of the form $\chi = \chi_1\chi_2...\chi_n$ and say which χ_i are principal. Our eigenvalue is the sum of subscripts of positions that weren't occupied by the principal character. If you want to know the multiplicity of p^i, you look at all of the eigenvalues and simply count how many times it exactly divides them.

Concluding remarks.

Generally speaking, when studying a graph, one technique is to encode the information into a matrix and then compute different algebraic and numeric properties of it. We can then use this information as displaying inherent properties of the graph itself.

We started by focusing our energy on the Hamming adjacency matrix and soon explored various other incidence matrices, including the n-cube and the Laplacian. By conjugating any adjacency matrix A_k with the group's character table, we were able to produce a diagonal matrix with eigenvalues $\sum_{e \in E_k} \chi(e)$. These eigenvalues can then be used to determine the elementary divisors of the matrix for all p such that $p \nmid |G|$. We made great progress and were able to compute the Smith normal form for different abelian Cayley graphs. It was a pleasure to work with Dr. Ducey.