The Combinatorics of Seidel Switching

Gregory Quenell

December 1997

1 Introduction – Graph spectra

Let $G = (V, E)$ be a finite, undirected graph with vertex set $\{v_1, v_2, \ldots, v_N\}$ and edge set $\{e_1, e_2, \ldots, e_M\}$. We allow G to have loops (edges connecting a vertex to itself) and parallel edges (two or more edges sharing the same endpoints). The *adjacency matrix* of G is the N-by-N integer matrix

$$A = (a_{ij})$$

in which a_{ij} is equal to the number of edges joining vertex v_i to vertex v_j (if $i \neq j$), and a_{ii} is equal to twice the number of loops at the vertex v_i.

We remark that A depends on the ordering of the vertices in the set V, but that a reordering of the vertices changes A only by permuting its rows and columns, so that any adjacency matrix for G is similar to A. Thus the set of eigenvalues of A, which we will call the *spectrum* of G, is independent of the ordering in the set V, and is completely determined by the graph structure of G.

The graph structure of G, on the other hand, is *not* in general determined by the spectrum of G. We have numerous examples of pairs of graphs which are isospectral but not isomorphic. (See [3] for several of them.) In this paper, we discuss a technique, called *Seidel switching*, for generating such pairs.

Seidel switching was introduced in [6], a study of point sets in elliptic spaces. In [3], there appears a very short algebraic proof of why graphs constructed via the Seidel technique are isospectral. Another algebraic proof is found in [5]. Here we offer a (somewhat longer) proof of the same fact, with a much more geometric flavor. This proof may fall into the category of mathematical folklore; it was shown to us by Robert Brooks ([1]).

2 Walks and the length spectrum

As before, let $G = (V, E)$ be an undirected graph, and let A be an adjacency matrix of G. Let $\lambda_1, \lambda_2, \ldots, \lambda_N$ be the eigenvalues of A. We remark that since A is symmetric, these eigenvalues are all guaranteed to be real.

The most basic connection between the spectrum of A and the geometric structure of G involves counting walks in G. In order to define and study walks, it will be convenient to view G as a directed graph. We do so by treating each undirected edge e joining vertices v and w as two directed edges: ε_1 leading from v to w, and ε_2 leading from w to v. Following this convention, an undirected loop at a vertex v in G will be treated as two directed loops at v. (We may think of these as leading around the undirected loop in opposite directions.)

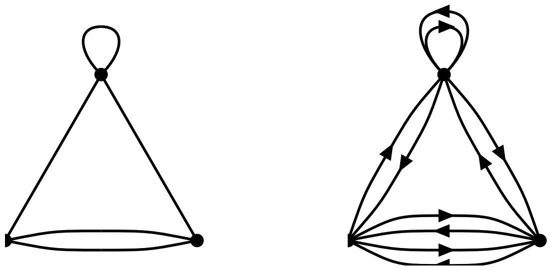

Undirected Directed

Figure 1: Directed version of an undirected graph

In Figure 1, we show an undirected graph beside a version of the same graph showing all the edges as pairs of directed edges.

For the rest of this paper, all of our graphs will be undirected, but the edge set E of each graph will most often be treated as a set of directed edges ε_i, with the understanding that for each ε_i in E, there is some opposite edge ε_j

also in E. We will use the notations $\mathrm{Init}(\varepsilon)$ and $\mathrm{Term}(\varepsilon)$ for the initial and terminal vertices, respectively, of the directed edge ε.

Let $G = (V, E)$ be a graph, and k a non-negative integer. A *k-walk* from vertex v to vertex w in G is a sequence of directed edges

$$\varepsilon_1 \, \varepsilon_2 \, \varepsilon_3 \, \cdots \, \varepsilon_k$$

in which $\mathrm{Init}(\varepsilon_1) = v$, $\mathrm{Term}(\varepsilon_k) = w$, and $\mathrm{Init}(\varepsilon_{i+1}) = \mathrm{Term}(\varepsilon_i)$ for $i = 1, \ldots, k-1$.

For a matrix B, let $[B]_{ij}$ denote the entry in the i^{th} row and j^{th} column of B. The following is well-known (see, for example, [4], p. 653).

Theorem 1 *Let $A = (a_{ij})$ be the adjacency matrix of an undirected graph G. For each integer $k \geq 0$, the number of k-walks from v_i to v_j in Γ is equal to*

$$[A^k]_{ij},$$

the ij^{th} entry of the k^{th} power of A.

A *closed k-walk* at v in G is a k-walk from v to v. It follows immediately from Theorem 1 that the number of closed k-walks at a vertex v_i in G is equal to $[A^k]_{ii}$, where A is the adjacency matrix of G. Taking the sum of these diagonal entries, we get the following.

Theorem 2 *Let G be a graph with adjacency matrix A, and let $\lambda_1, \lambda_2, \ldots, \lambda_N$ be the eigenvalues of A. For each integer $k \geq 0$, the total number of closed k-walks in G is equal to*

$$\lambda_1^k + \lambda_2^k + \cdots + \lambda_N^k.$$

Proof The total number of closed k-walks in G is equal to the sum over all the vertices v_i in G of the number of closed k-walks at v_i. By the observation above, this is equal to

$$\sum_{i=1}^{N} [A^k]_{ii} \quad = \quad \mathrm{Tr}(A^k).$$

Now if λ is an eigenvalue of A, then λ^k is an eigenvalue of A^k (with the same eigenvector), so the eigenvalues of A^k are

$$\lambda_1^k,\ \lambda_2^k,\ \ldots,\ \lambda_N^k,$$

and the trace of A^k is

$$\lambda_1^k + \lambda_2^k + \cdots + \lambda_N^k,$$

as required. ■

For each integer $k \geq 0$, let $\ell_k(G)$ denote the number of closed k-walks in the graph G. The *length spectrum* of G is the sequence

$$\ell_0(G),\ \ell_1(G),\ \ell_2(G), \ldots.$$

Theorem 2 says that

$$\ell_k(G) \quad = \quad \sum_{i=1}^{N} \lambda_i^k,$$

where the numbers λ_i are the eigenvalues of A, and thus that the length spectrum of G is determined by the spectrum of G.

It is also true that the spectrum of G is determined by the length spectrum of G. In fact, the spectrum of G, the complete set of eigenvalues of A, is determined by the numbers

$$\ell_1(G),\ \ell_2(G),\ \ldots,\ \ell_N(G)$$

through a set of relations called *Newton's formulas*. Briefly (see [2]), given complex numbers $r_1, r_2, \ldots, r_n$, the polynomial

$$p(x) \quad = \quad (x - r_1)(x - r_2) \cdots (x - r_n)$$

may be expanded as

$$p(x) \quad = \quad \sigma_0 x^n - \sigma_1 x^{n-1} + \cdots + (-1)^i \sigma_i x^{n-i} + \cdots + (-1)^n \sigma_n,$$

where σ_i is the i^{th}-degree elementary symmetric function in $r_1, r_2, \ldots, r_n$, and we have set $\sigma_0 = 1$ for convenience. For each non-negative integer i let

$$\pi_i \quad = \quad r_1^i + r_2^i + \cdots + r_n^i$$

be the i^{th} moment of the roots of $p(x)$. Then for k between 0 and n (inclusive), we have

$$\sigma_k \;=\; \frac{1}{k}\left[\sum_{i=1}^{k}(-1)^{i+1}\sigma_{k-i}\pi_i\right].$$

Applying this relation recursively, we can recover σ_1 through σ_n (and thus the polynomial $p(x)$) from the numbers π_1 through π_n. In the present instance, this tells us that we can reconstruct the characteristic polynomial of an N-vertex graph G from the numbers $\ell_1(G), \ell_2(G), \ldots, \ell_N(G)$.

If two graphs have a common length spectrum, they are said to be *length isospectral*. By the argument above, we have

Corollary 3 *Two graphs are isospectral if and only if they are length isospectral.*

3 Seidel switching

Here is the recipe for producing a pair of isospectral graphs using Seidel switching.

We begin with two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$. The graph G_2 must contain an even number of vertices and must be regular (that is, every vertex must have the same number of edges incident upon it). There are no restrictions on G_1.

We introduce a set $\mathcal{E}$ of (undirected) edges which join each vertex in G_1 to exactly half the vertices in G_2. We may choose the edges in $\mathcal{E}$ in any way we like, as long as there are exactly $\frac{|V_2|}{2}$ edges in $\mathcal{E}$ incident on each vertex in G_1, and the other ends of these edges are incident on some $\frac{|V_2|}{2}$ different vertices in G_2.

We next form the set $\mathcal{E}^C$ of undirected edges, which is determined by $\mathcal{E}$ as follows: for each vertex v in G_1 and each vertex w in G_2, there is an edge between v and w in $\mathcal{E}^C$ if and only if there is no edge between v and w in $\mathcal{E}$. That is, the edges in $\mathcal{E}^C$ connect each vertex in G_1 to the "other half" of the vertices in G_2.

We now form two graphs,

$$\Gamma_A = (V_1 \cup V_2, E_1 \cup E_2 \cup \mathcal{E}) \tag{1}$$

and

$$\Gamma_B = (V_1 \cup V_2, E_1 \cup E_2 \cup \mathcal{E}^C). \tag{2}$$

In the next section, we will prove that Γ_A and Γ_B thus formed are always length isospectral, and thus isospectral.

First, we illustrate this procedure with an example. We begin with the two graphs G_1 and G_2 in Figure 2. Note that G_2 has an even number of vertices and is 2-regular.

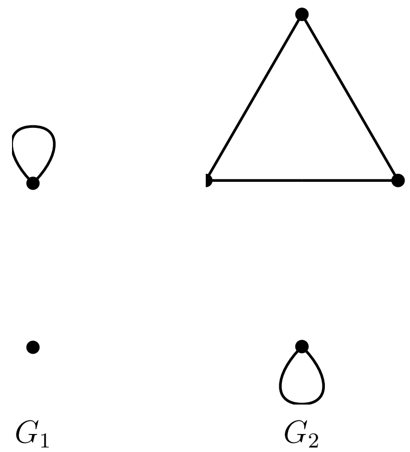

Figure 2: Graphs G_1 and G_2 for the Seidel construction

In Figure 3, we introduce the set $\mathcal{E}$, shown with dotted lines, connecting each vertex of G_1 to two vertices of G_2. The set $\mathcal{E}^C$ is shown with shaded lines.

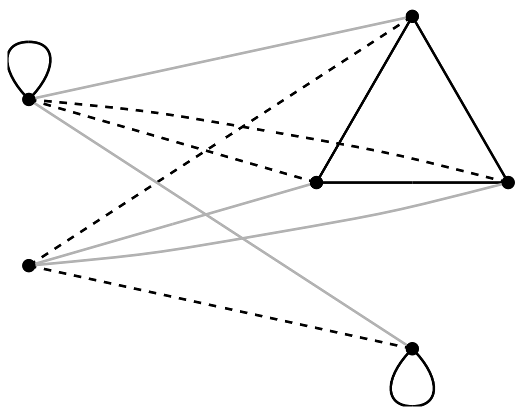

Figure 3: The graphs G_1 and G_2, with $\mathcal{E}$ and $\mathcal{E}^C$

The graphs $\Gamma_A = (V_1 \cup V_2, E_1 \cup E_2 \cup \mathcal{E})$ and $\Gamma_B = (V_1 \cup V_2, E_1 \cup E_2 \cup \mathcal{E}^C)$ are shown in Figures 4 and 5 respectively.

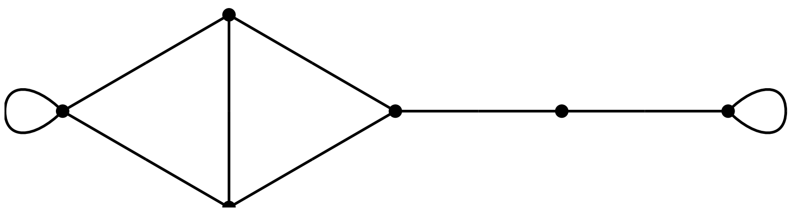

Figure 4: The graph Γ_A

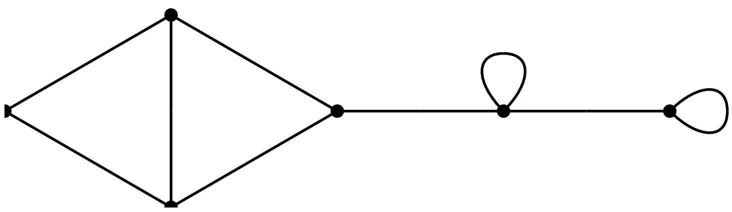

Figure 5: The graph Γ_B

The adjacency matrices of Γ_A and Γ_B share the characteristic polynomial

$$X^6 - 4X^5 - 3X^4 + 18X^3 + 4X^2 - 16X + 4,$$

so the graphs are indeed isospectral.

The construction of the isospectral pair Γ_A and Γ_B (lines (1) and (2)) is determined once G_1, G_2, and $\mathcal{E}$ have been specified, so we will refer to the entire construction above (involving G_1, G_2, $\mathcal{E}$, $\mathcal{E}^C$, Γ_A, and Γ_B) as the *Seidel pair* $(G_1, G_2, \mathcal{E})$. We sympathize with the reader who finds it disturbing to have a "pair" with three elements.

4 A combinatorial proof

Suppose we are given a Seidel pair $(G_1, G_2, \mathcal{E})$.

Given vertices v and w in V_1 (the vertex set of G_1) and a non-negative integer k, we define a (v, w, k) *A-patch* to be a sequence of (directed) edges

$$\varepsilon_0\, \varepsilon_1\, \varepsilon_2 \cdots \varepsilon_k\, \varepsilon_{k+1}$$

in which

1. $\varepsilon_0 \in \mathcal{E}$ and $\varepsilon_{k+1} \in \mathcal{E}$
2. $\mathrm{Init}(\varepsilon_0) = v$ and $\mathrm{Term}(\varepsilon_{k+1}) = w$
3. $\varepsilon_i \in E_2$ for $i = 1, \ldots, k$
4. $\mathrm{Term}(\varepsilon_i) = \mathrm{Init}(\varepsilon_{i+1})$ for $i = 0, \ldots, k$.

That is, a (v, w, k) A-patch is a walk in Γ_A from v to w in which the first and last edges, and only the first and last edges, are in $\mathcal{E}$. Similarly, a (v, w, k) *B-patch* is a walk from v to w in Γ_B in which just the first and last edges are in $\mathcal{E}^C$. Formally, a (v, w, k) B-patch is a sequence of directed edges

$$\varepsilon_0\, \varepsilon_1\, \varepsilon_2 \cdots \varepsilon_k\, \varepsilon_{k+1}$$

satisfying

1. $\varepsilon_0 \in \mathcal{E}^C$ and $\varepsilon_{k+1} \in \mathcal{E}^C$
2. $\mathrm{Init}(\varepsilon_0) = v$ and $\mathrm{Term}(\varepsilon_{k+1}) = w$
3. $\varepsilon_i \in E_2$ for $i = 1, \ldots, k$
4. $\mathrm{Term}(\varepsilon_i) = \mathrm{Init}(\varepsilon_{i+1})$ for $i = 0, \ldots, k$.

The core of the combinatorial proof of the validity of Seidel switching is the following lemma about A-patches and B-patches.

Lemma 4 *Given a Seidel pair $(G_1, G_2, \mathcal{E})$, for each pair v and w of vertices in V_1 and each non-negative integer k, the number of (v, w, k) A-patches is equal to the number of (v, w, k) B-patches.*

In the course of proving Lemma 4, we will need the following simple result.

Lemma 5 *Let G be an r-regular graph. Given a vertex v in G and a non-negative integer k, the number of k-walks in G which begin at v is r^k. Also, the number of k-walks in G which end at v is r^k.*

Proof The number of 0-walks beginning at v is certainly 1, and the number of 1-walks beginning at v is r, because there are r edges originating at v. Suppose we have a $(k-1)$-walk beginning at v and ending at some vertex w. We can extend this to a k-walk in exactly r ways, since there are r edges originating at w. Thus the number of k-walks beginning at v is r times the number of $(k-1)$-walks beginning at v. By induction, this number is

$$r \cdot r^{k-1} \;=\; r^k.$$

The proof of the second part of the lemma is entirely analogous, and we omit it. ■

We are now ready to undertake the proof of Lemma 4.

Proof Fix an integer $k \geq 0$ and vertices v and w in V_1. The recipe for a Seidel pair guarantees that each vertex in V_2 is joined to v by an edge in $\mathcal{E}$ or by an edge in $\mathcal{E}^C$, but not both. That is, each vertex in V_2 is adjacent to v in Γ_A or in Γ_B, but not both. Thus we can partition the vertices in V_2 as

$$V_2 \;=\; V_{(v,A)} \cup V_{(v,B)}$$

where $V_{(v,A)}$ is the set of vertices in V_2 which are adjacent to v in Γ_A and $V_{(v,B)}$ is the set of those which are adjacent to v in Γ_B. For the same reason, we have a second partition

$$V_2 \;=\; V_{(w,A)} \cup V_{(w,B)}$$

where the vertices in $V_{(w,A)}$ are adjacent to w in Γ_A and those in $V_{(w,B)}$ are adjacent to w in Γ_B.

Furthermore, since the edges in $\mathcal{E}$ join each of v and w to exactly half of the vertices in V_2, we know that

$$|V_{(v,A)}| = |V_{(v,B)}| = |V_{(w,A)}| = |V_{(w,B)}| = \frac{|V_2|}{2}.$$

We now partition all the k-walks in G_2 into four sets,

$$W_{AA} \cup W_{AB} \cup W_{BA} \cup W_{BB}$$

according to where their beginning and ending vertices fall. Consider the table

		Location of starting vertex	
		$V_{(v,A)}$	$V_{(v,B)}$
Location of ending vertex	$V_{(w,A)}$	W_{AA}	W_{AB}
	$V_{(w,B)}$	W_{BA}	W_{BB}

The set W_{AA} contains all the k-walks in G_2 which begin at a vertex in $V_{(v,A)}$ and end at a vertex in $V_{(w,A)}$. The set W_{AB} contains all the k-walks in G_2 which begin at a vertex in $V_{(v,A)}$ and end at a vertex in $V_{(w,B)}$. Since

$$V_{(w,A)} \cup V_{(w,B)}$$

is a partition of V_2, the set $W_{AA} \cup W_{AB}$ comprises all the k-walks in G_2 which begin at a vertex in $V_{(v,A)}$ and end anywhere in V_2. Since G_2 is r-regular and there are exactly $\frac{|V_2|}{2}$ vertices in $V_{(v,A)}$, we apply Lemma 5 to get

$$|W_{AA}| + |W_{AB}| = |V_{(v,A)}|r^k = \frac{|V_2|}{2}r^k. \tag{3}$$

Similarly, $W_{AB} \cup W_{BB}$ comprises all the k-walks in G_2 which begin anywhere in V_2 and end at a vertex in $V_{(w,B)}$, and so

$$|W_{AB}| + |W_{BB}| = |V_{(w,B)}|r^k = \frac{|V_2|}{2}r^k. \tag{4}$$

From (3) and (4), we get

$$|W_{AA}| + |W_{AB}| \;=\; |W_{AB}| + |W_{BB}|,$$

from which it follows that

$$|W_{AA}| \;=\; |W_{BB}|.$$

To complete the proof, we observe that each k-walk in W_{AA} corresponds to exactly one (v, w, k) A-patch, and each k-walk in W_{BB} corresponds to exactly one (v, w, k) B-patch. Explicitly, the correspondence (in the W_{AA} case) associates the k-walk

$$\varepsilon_1\, \varepsilon_2\, \cdots\, \varepsilon_k$$

with the patch

$$\varepsilon_0\, \varepsilon_1\, \varepsilon_2\, \cdots\, \varepsilon_k\, \varepsilon_{k+1},$$

where ε_0 is the edge in $\mathcal{E}$ joining v to the initial vertex of ε_1 (suitably directed) and ε_{k+1} is the edge in $\mathcal{E}$ joining the terminal vertex of ε_k to w (again, suitably directed). The correspondence in the W_{BB} case is constructed analogously. ■

Since the sets W_{AA} and W_{BB} have the same cardinality, we have the following.

Corollary 6 *Let v and w be vertices in V_1 and k a non-negative integer. Let $\mathcal{P}^A_{(v,w,k)}$ be the set of (v, w, k) A-patches from v to w and $\mathcal{P}^B_{(v,w,k)}$ be the set of (v, w, k) B-patches from v to w. Then there is a one-to-one correspondence*

$$\mathcal{P}^A_{(v,w,k)} \;\leftrightarrow\; \mathcal{P}^B_{(v,w,k)}.$$

We are now ready to show that the graphs Γ_A and Γ_B in a Seidel pair are length isospectral.

Theorem 7 *Let k be a non-negative integer and $(G_1, G_2, \mathcal{E})$ a Seidel pair. There is a one-to-one correspondence between the set of all closed k-walks in Γ_A and the set of all closed k-walks in Γ_B.*

Proof Every closed k-walk W in Γ_A is of one of three types:

$$\begin{array}{rl} \text{Type I:} & W \text{ is contained in } G_1 \\ \text{Type II:} & W \text{ is contained in } G_2 \\ \text{Type III:} & W \text{ contains some edges in } \mathcal{E}. \end{array}$$

Similarly, if Z is a closed k-walk in Γ_B, then Z is of one of these three types:

$$\begin{array}{rl} \text{Type I}': & Z \text{ is contained in } G_1 \\ \text{Type II}': & Z \text{ is contained in } G_2 \\ \text{Type III}': & Z \text{ contains some edges in } \mathcal{E}^C. \end{array}$$

Given a closed k-walk W in Γ_A, if W is of Type I or Type II, then W is in fact a closed k-walk in Γ_B (of Type I$'$ or Type II$'$), so we may associate such a W to a corresponding closed k-walk in Γ_B using the identity mapping.

To complete the proof, we need to construct a bijection from the set of Type III closed k-walks in Γ_A to the set of Type III$'$ closed k-walks in Γ_B. Let

$$W \quad = \quad \varepsilon_1 \, \varepsilon_2 \, \cdots \, \varepsilon_k$$

be a closed k-walk of Type III in Γ_A. Because W is closed and contains some edge in $\mathcal{E}$, it must contain some first edge ε_m which is in $\mathcal{E}$ and leads from a V_1 vertex to a V_2 vertex. We permute the edges in W cyclically so that ε_m comes first, and consider the resulting walk,

$$\overline{W} \quad = \quad \varepsilon_m \, \varepsilon_{m+1} \, \varepsilon_{m+2} \, \cdots \, \varepsilon_{m+k-1}$$

where the subscripts are read modulo k. This is a closed k-walk, beginning and ending at some vertex $v = \mathrm{Init}(\varepsilon_m)$ in V_1. As such, it may be viewed as a sequence

$$\overline{W} \quad = \quad P_1 \, W_1 \, P_2 \, W_2 \, \cdots \, P_j \, W_j \tag{5}$$

where each P_i is a (v_i, w_i, l_i) A-patch and each W_i is a walk, contained entirely in G_1, from w_i to v_{i+1}, reading these subscripts modulo j.

By Corollary 6, for each patch P_i, there exists a corresponding (v_i, w_i, l_i) B-patch, which we will denote P_i'. We replace each P_i in $\overline{W}$ by its corresponding P_i' to get a new walk

$$\overline{W}' \;=\; P_1' \, W_1 \, P_2' \, W_2 \, \cdots \, P_j' \, W_j.$$

This is again a closed k-walk at $v = \mathrm{Init}(\varepsilon_m)$, but each edge ε_h of $\overline{W}$ that was contained in $\mathcal{E}$ has been replaced by an edge ε_h' in $\mathcal{E}^C$. Thus $\overline{W}'$ is a closed k-walk in Γ_B. To complete the mapping, we write

$$\overline{W}' \;=\; \varepsilon_m' \, \varepsilon_{m+1}' \, \varepsilon_{m+2}' \, \cdots \, \varepsilon_{m+k-1}'$$

where each ε_h' is either equal to ε_h (if ε_h was in some W_i in (5)) or is the edge corresponding to ε_h in some P_i'. Finally, we apply another cyclic permutation to the edges in $\overline{W}'$ to recover

$$W' \;=\; \varepsilon_1' \, \varepsilon_2' \, \cdots \, \varepsilon_k',$$

which is indeed a Type III$'$ closed k-walk in Γ_B.

To invert this mapping $W \mapsto W'$, we proceed as follows. A closed k-walk W' in Γ_B contains some first edge ε_m' in $\mathcal{E}^C$ which leads from V_1 to V_2. We apply a cyclic permutation to the edges in W' so that ε_m' comes first, and then replace each (v_i, w_i, l_i) B-patch in the resulting closed k-walk with the corresponding (v_i, w_i, l_i) A-patch, whose existence is guaranteed by Corollary 6. Finally we undo our cyclic permutation so that the image of ε_1' comes first.

To see that this procedure really does invert our original mapping, we need only observe that the first edge in W leading from V_1 to V_2 is in exactly the same position as the first edge in W' leading from V_1 to V_2 and that replacing a B-patch with an A-patch is the inverse operation to replacing an A-patch with a B-patch. ■

5 Seidel switching and regularity

The spectrum of a graph determines some of the graph's geometry, but the existence of Seidel pairs (and the ease with which they can be constructed) shows that the spectrum of a graph is far from determining its isomorphism type. If we introduce a new restriction and consider only regular graphs, then the spectrum seems to become somewhat stronger, in the sense that isospectral pairs of regular graphs appear to be rarer than isospectral pairs among non-regular graphs.

Thus it is of interest to construct Seidel pairs of regular graphs. The following theorem gives some necessary conditions for a Seidel pair to be regular.

Theorem 8 *Suppose $(G_1, G_2, \mathcal{E})$ is a Seidel pair, and that Γ_A is q-regular. Then*

1. $|V_1|$ *is even.*
2. G_1 *is regular.*
3. $\frac{|V_1|}{2} + r = \frac{|V_2|}{2} + s = q$, *where r is the valency of G_2 and s is the valency of G_1.*
4. Γ_B *is q-regular.*

Proof G_2 is r-regular to begin with, so in Γ_A, each vertex in V_2 must be an endpoint of the same number of edges in $\mathcal{E}$, and this number must be $\frac{|\mathcal{E}|}{|V_2|}$. Now by the Seidel recipe, we know that

$$|\mathcal{E}| \quad = \quad \frac{|V_1| \cdot |V_2|}{2},$$

so the number of edges in $\mathcal{E}$ incident on each vertex in G_2 must be $\frac{|V_1|}{2}$, showing that $|V_1|$ is even, and that

$$q \quad = \quad r + \frac{|V_1|}{2}. \tag{6}$$

Let v be a vertex in V_1, and let s be the valency of v in G_1. In the graph Γ_A, there are $\frac{|V_2|}{2}$ new edges incident on v, so the valency of v as a vertex in Γ_A is

$$s + \frac{|V_2|}{2}.$$

But Γ_A is q-regular, so we get

$$s = q - \frac{|V_2|}{2},$$

independent of the choice of v. Thus G_1 is s regular, and

$$q \;=\; s + \frac{|V_2|}{2}. \tag{7}$$

The third assertion has been proved in lines (6) and (7).

Since $\mathcal{E}^C$ contains edges joining each vertex in V_1 to exactly half the vertices in V_2, we know that the valency of each vertex in V_1 in the graph Γ_B is

$$s + \frac{|V_2|}{2} \;=\; q.$$

Meanwhile, if v is a vertex in V_2, then there are exactly $\frac{|V_1|}{2}$ vertices in V_1 which are joined to v by edges in $\mathcal{E}$. This means that there are $\frac{|V_1|}{2}$ vertices in V_1 which are *not* joined to v by edges in $\mathcal{E}$, it is exactly this set of vertices which must be joined to v by edges in $\mathcal{E}^C$. Thus the valency of v in the graph Γ_B is

$$r + \frac{|V_1|}{2} \;=\; q.$$

■

We note that conditions 1, 2, and 3 are not sufficient for the regularity of Γ_A and Γ_B, but they do tell us where to look for examples of regular Seidel pairs.

6 Small, regular Seidel pairs

In particular, we can use Theorem 8 to look for the smallest examples of regular Seidel pairs in which Γ_A and Γ_B are not isomorphic. We have

Corollary 9 *If $(G_1, G_2, \mathcal{E})$ is a Seidel pair in which either G_1 or G_2 contains only two vertices, then Γ_A is isomorphic to Γ_B.*

Proof Suppose G_1 contains only two vertices v_1 and v_2. Since G_1 is regular, there exists an automorphism of G_1 which interchanges v_1 and v_2. Because $\mathcal{E}$ joins v_1 to half the vertices in G_2 and v_2 to the other half, a vertex w in G_2 is adjacent to v_1 in Γ_A if and only if w is not adjacent to v_2 in Γ_A. But w is not adjacent to v_2 in Γ_A if and only if w *is* adjacent to v_2 in Γ_B. Thus

$$w \sim v_1 \text{ in } \Gamma_A \quad \Leftrightarrow \quad w \sim v_2 \text{ in } \Gamma_B.$$

This, together with the fact that interchanging v_1 and v_2 is an automorphism of G_1, implies that interchanging v_1 and v_2 is an isomorphism taking Γ_A to Γ_B.

An analogous argument applies if G_2 has only two vertices. ■

With a little more work, it can be shown that if $|V_1| = |V_2| = 4$ and G_1 is 1-regular, then the graphs Γ_A and Γ_B in any Seidel pair $(G_1, G_2, \mathcal{E})$ are isomorphic. (The idea is that G_1 must be the union of two K_2s, and the automorphism group of this graph is large enough to induce an isomorphism between Γ_A and Γ_B for any choice of $\mathcal{E}$.)

There are at least three regular Seidel pairs (with Γ_A not isomorphic to Γ_B) in which $|V_1| = |V_2| = 4$ and G_1 and G_2 are 2-regular. We present one of them, identified in [1]. Figure 6 shows the four-vertex, 2-regular graphs G_1 (on the left) and G_2 (on the right). The edges in $\mathcal{E}$ are shown in grey.

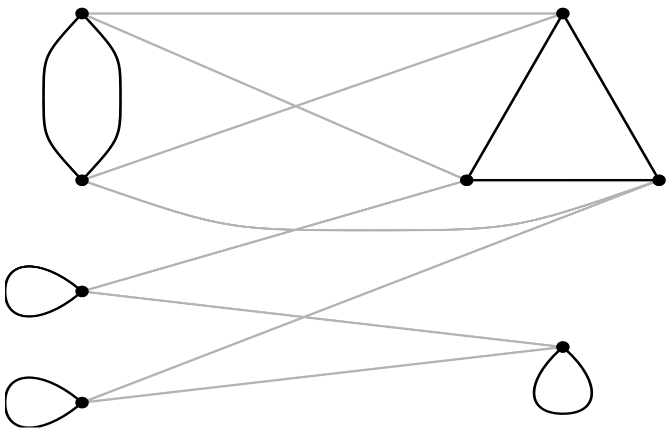

Figure 6: G_1, $\mathcal{E}$, and G_2

The graphs Γ_A and Γ_B in this Seidel pair (shown in Figure 7) are pleasingly symmetric and clearly non-isomorphic.

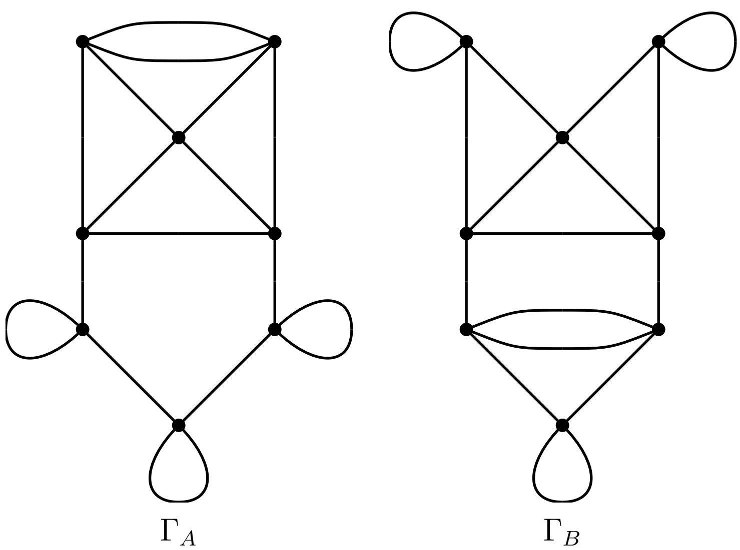

Figure 7: A 4-regular Seidel pair with $|V_1| = |V_2| = 4$

Looking further, Theorem 8 and Corollary 9 do not rule out the possibility of an interesting regular Seidel pair with $|V_1| = 4$, $|V_2| = 6$, $r = 1$, and $s = 2$. In fact, such a pair exists, and seems to be the smallest example of a Seidel pair of regular graphs containing no loops or parallel edges. We let G_1 be the union of two K_2s and G_2 be a six-vertex cycle graph. Figure 8 shows G_1, G_2 and the edges in $\mathcal{E}$.

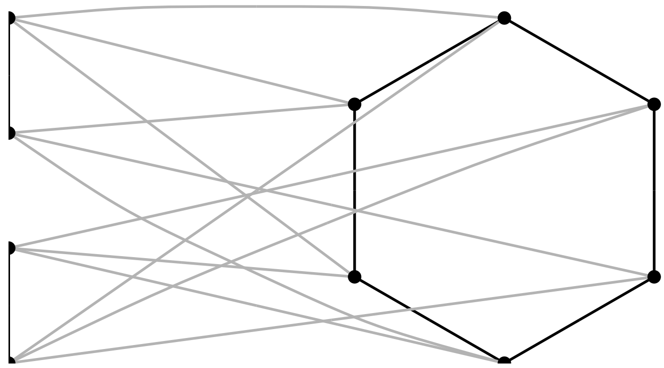

Figure 8: G_1, $\mathcal{E}$, and G_2

Again, the graphs Γ_A and Γ_B arising from this pair exhibit pleasing symmetries. We show them in Figures 9 and 10. The circled vertices in each of these figures are to be identified.

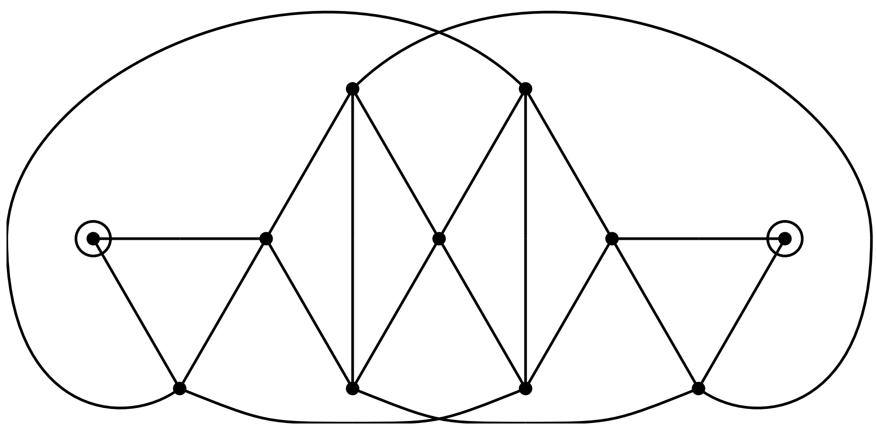

Figure 9: Γ_A of a 4-regular, ten-vertex Seidel pair

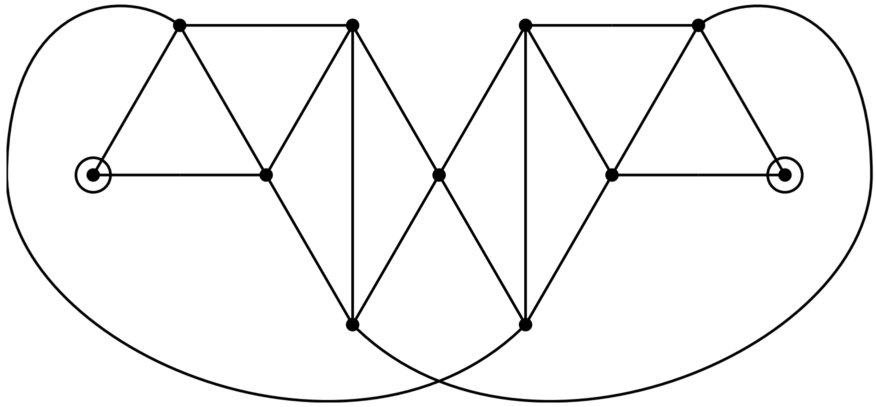

Figure 10: Γ_B of a 4-regular, ten-vertex Seidel pair

References

[1] Robert Brooks. Seidel switching and isospectral graphs. Lecture at Dartmouth College, February 1996.

[2] Allan Clark. *Elements of Abstract Algebra.* Dover Publications, Mineola, NY, 1984.

[3] Dragoš M. Cvetković, Michael Doob, and Horst Sachs. *Spectra of Graphs.* Academic Press, New York, NY, 1979.

[4] Susanna S. Epp. *Discrete Mathematics with Applications.* PWS Publishing Company, Boston, second edition, 1995.

[5] C. D. Godsil and B. D. McKay. Constructing cospectral graphs. *Aequationes Mathematicae*, 25:257–268, 1982.

[6] J. H. van Lint and J. J. Seidel. Equilateral point sets in elliptic geometry. *Proceedings. Koninklijke Nederlanse Academie van Wetenschppen, Series A*, 69:335–348, 1966.