PETERSEN SIGNED GRAPHS
PARTIAL DRAFT June 8, 2009

THOMAS ZASLAVSKY
BINGHAMTON UNIVERSITY OF SUNY

Contents

1. The Beginning

Deservedly a favorite in graph theory, the Petersen graph P illustrates many of the important properties of graphs, either as a non-trivial example or, remarkably often, as a counterexample. For instance, it is cubic and has large girth for its size but it is neither vertex nor edge 3-colorable. It is very highly symmetric and strongly regular but not Hamiltonian. It is nonplanar, which is obvious from Kuratowski's theorem, but its crossing number is greater than 1. It is so interesting that it has been the subject of a book [15], which uses it as a springboard to introduce many areas of graph theory. And it distinguishes the flag of the CombinaTexas conference [8].

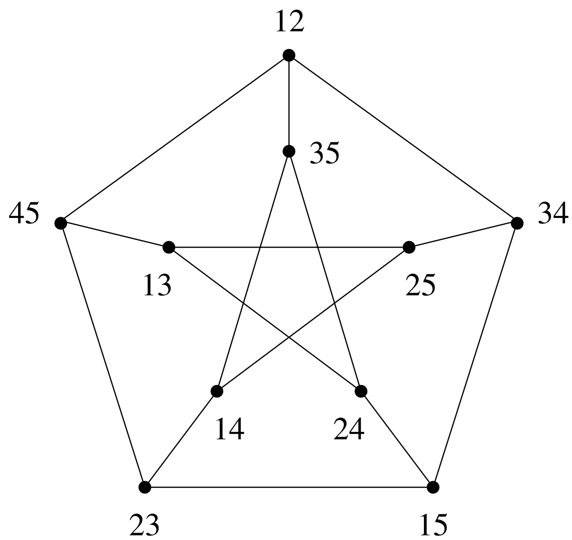

Figure 1.1. The Petersen graph. (Traditional.) The vertex labels are derived from its representation as $\overline{L(K_5)}$.

The Petersen graph is also a beautiful example in the theory of signed graphs. A *signed graph* is a graph with signed edges; that is, each edge is either positive or negative. A simple graph Γ can be treated as a signed graph in four natural ways. There are $+\Gamma$, where all edges are positive; $-\Gamma$, where all edges are negative; K_Γ, the complete graph of the same order whose negative subgraph is Γ and whose positive subgraph is the complement $\bar{\Gamma}$; and of course $K_{\bar{\Gamma}}$, which is the negative of K_Γ. Each of these is interesting for different reasons. Usually, $+\Gamma$ behaves just like the unsigned graph Γ. $-\Gamma$, on the contrary, is quite a different animal. K_Γ seems to live in another universe; still, it does have certain subtle connections with Γ. Yet other signatures arise from nonorientable surface embeddings. These constructions applied to the Petersen graph produce signed graphs that exemplify a great many of the significant and interesting properties of and questions about signed graphs. I will survey some of them here, and even produce a few new general (though small) results.

Part I. Graph Theory

2. Sets, Graphs, and Signed Graphs: Definitions

We'll need a few definitions right away.

The 'canonical' set of n elements is $[n] := \{1, 2, \dots, n\}$. The power set of E is $\mathcal{P}(E)$. The class of k-element subsets of E is $\mathcal{P}_k(E)$. The class of all partitions of a set V is Π_V. The *trivial partition* of a set has just one block. A partition of a graph is simply a partition of its vertex set.

Some notation I use more or less frequently:

* $\Gamma := (V, E)$ and $n := |V|$, the order of Γ; sometimes I will number the vertices $V = \{v_1, \dots, v_n\}$, but often they will be numbered differently, as is convenient.
* The number of connected components is $c(\Gamma)$. $\pi(\Gamma)$ denotes the partition of V into the vertex sets of the components of Γ; this partition has $c(\Gamma)$ blocks.
* The union of graphs, $\Gamma_1 \cup \Gamma_2$, is not disjoint unless their vertex sets are disjoint. We may unite two graphs on the same or overlapping vertex sets. Disjoint union is $\Gamma_1 \mathbin{\dot\cup} \Gamma_2$.
* The (open) neighborhood of a vertex u in Γ is $N(u) := N_\Gamma(u)$. The closed neighborhood is $N[u] := N(u) \cup \{u\}$. The open and closed neighborhoods of a set $X \subseteq V$ are $N(X) := \big(\bigcup_{u \in X} N(u)\big) \setminus X$ and $N[X] := \bigcup_{u \in X} N[u]$.
* A *circle* is a 2-regular connected graph, or its edge set.
* A path of length l is P_l.
* A *cut* is a nonvoid edge set $E(X, X^c)$ that consists of all edges between a vertex subset $X \subseteq V$ and its complement X^c.
* A *vertex star* is one kind of cut: it is the set of all edges incident to a vertex.
* A *walk* is a sequence of vertices and edges, $v_0 e_1 v_1 e_2 \cdots v_{l-1} e_l v_l$, such that e_i has endpoints v_{i-1}, v_i. A *path* is a walk without repeated edges or vertices.
* A graph is *k-regular* if every vertex has degree k. If there are loops, they count double in the degree.
* The distance between two edges is the number of vertices in a shortest path that joins a vertex of one edge to a vertex of the other edge. For instance, adjacent edges have distance 1. This is the same as distance in the line graph.
* The contraction of Γ by an edge set S is written Γ/S.

A *signed graph* is a pair $\Sigma = (|\Sigma|, \sigma)$ consisting of an *underlying graph* $|\Sigma| = (V, E)$ and a signature $\sigma : E \to \{+1, -1\}$. For the most part we think of $\{+1, -1\}$ as a multiplicative group, not as numbers. Again, the order of the graph is $n := |V|$. The graph need not be simple, but if it is, I call Σ a *signed simple graph*. Do not confuse this with a *simply signed graph*, which is a signed graph in which there are no parallel edges with the same sign; a simply signed graph may have vertices joined simultaneously by a positive and a negative edge, and it may have negative (but not positive) loops. I will often assume that all graphs are without loops; that makes many things simpler than they would otherwise be.

In a signed graph, E^+ and E^- are the sets of positive and negative edges. Σ^+ and Σ^- are the spanning subgraphs (unsigned) whose edges sets are E^+ and E^-, respectively.

Signed graphs Σ and Σ' are *isomorphic*, $\Sigma \cong \Sigma'$, when there is a sign-preserving isomorphism of underlying graphs.

The number of connected components of Σ that are balanced is $b(\Sigma)$.

3. The Petersen Graph

I like to think that what makes the Petersen graph so good, and certainly what gives it its high symmetry, is that it has a simple and elegant technical description. In $\mathcal{P}_2([5])$, write a subset of [5] as ij, without set notation. Let us call ij and kl *adjacent* when they are disjoint. The graph with $V = \mathcal{P}_2([5])$ and the disjointness adjacency relation is the Petersen graph P. (This is not the original definition; see [15, Section 9.7].)

With this definition it's obvious that P is trivalent, is edge transitive, and has the entire symmetric group of degree 5 as an automorphism group (and note that this is the whole automorphism group). Some other properties are not as easy to detect; for instance, the chromatic number $\chi(P) = 3$ and the chromatic index $\chi'(P) = 4$, and non-Hamiltonicity.

Another way to state the definition is that $P = \overline{L(K_5)}$, the complement of the line graph of K_5; thus we may write $-K_P = K_{\bar{P}} = K_{L(K_5)}$.

The Petersen graph has lovely combinatorial properties that will be very useful to us.

Lemma 3.1. *For each edge e in P there are exactly two edges at distance 3, and they are at distance 3 from each other. All other edges have distance at most 2 from e. Furthermore, any two sets of three edges at mutual distances 3 are equivalent by an automorphism of P.*

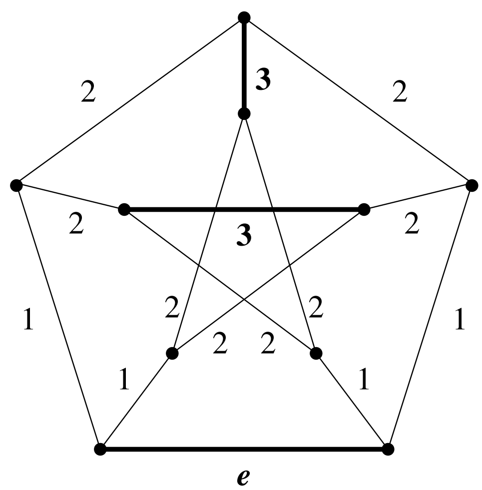

FIGURE 3.1. The distances from e of edges in the Petersen graph.

Proof. Specifically, the edges at distance 3 from $e = \{ij, kl\}$ are $\{ik, jl\}$ and $\{il, jk\}$.

For a pictorial proof, due to edge transitivity we need to consider only one edge e. The lemma is visible in Figure 3.1: the two edges at distance 3 from e are unique. Any automorphism that carries e to an edge f also carries the two distance-3 edges of e to those of f. □

Lemma 3.2. *A vertex set in P has all or none of the following properties.*

(i) *It consists of the four vertices labelled ij for a fixed i and all $j \neq i$.*
(ii) *It is a maximum independent set.*
(iii) *It is $N(u, v)$ for an edge uv.*
(iv) *It consists of the four vertices that are not covered by three edges at mutual distance 3.*

Proof. We rely on the edge and vertex transitivity of P to make the proof pictorial.

Since we know the clique structure of a line graph—a clique comes from a triangle or a vertex star in the base graph—this gives us the exact independence number $\alpha(P) = 4$ and identifies all the maximum-sized independent vertex sets: they have the form $V(i) := \{ij : j \neq i\}$ for some fixed i as in (i).

By inspection, the vertices at distance 1 from edge $\{12, 34\}$, for instance, are those of the form $5j$, again as in (i).

Also by inspection, a set $V(5)$, for example, is precisely the set of vertices not covered by the edges $\{12, 34\}, \{13, 24\}, \{14, 23\}$. □

Signatures of P.

There are many ways to associate a signed graph to the Petersen graph, but I will focus my attention on four, or rather seven since they come in pairs: $+P$ and $-P$ (which are actually quite different), K_P and $-K_P$ (which are not very different), the signature P_D shown in Figure 3.2 and its negative $-P_D$ (which are also less different than they appear; see Figure 4.2), and $P(e)$, which has just one negative edge e. Plainly, P has many other signatures, but only six are essential; Section 8 lists them.

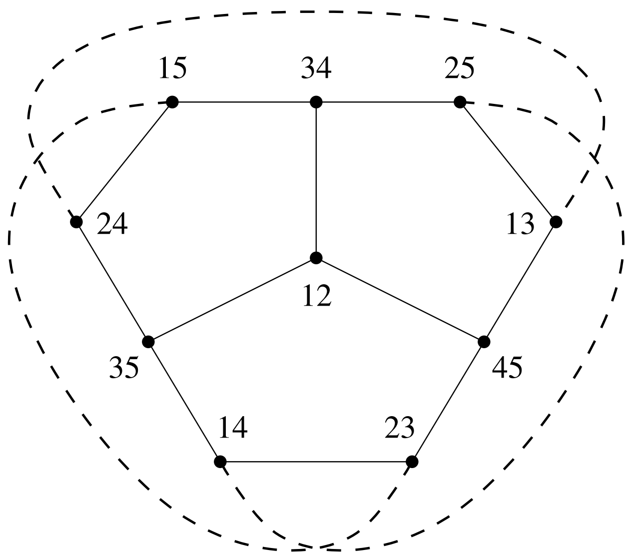

FIGURE 3.2. The Petersen graph drawn with three-fold symmetry around a central vertex. The dashed edges are negative in the signed graph P_D; they are alternate edges of the "special hexagon" $15, 24, 13, 25, 14, 23, 15$.

Let's count the positive and negative short circles in the various signed graphs. The numbers for signed Petersens are easy; bear in mind that in $-P$, as compared with $+P$, the circles of odd length change sign while those of even length do not.

We want to count negative triangles in K_P. A triangle is negative when it contains precisely one negative edge. There are 15 negative edges and for each one there are 4 vertices that are positively adjacent to both endpoints; that makes 60 negative triangles.

We also want the number of negative quadrilaterals. Such a quadrilateral could have either one or three negative edges. A negative quadrilateral with one negative edge is a path of length 3 with a negative edge in the middle, together with a last edge that isn't negative. To construct this we choose a Petersen edge e; that leaves a set X of 4 vertices that are positively adjacent to both endpoints of e. We need two of them to complete the quadrilateral, but we need two that do not support a Petersen edge. Since X supports only edges at distance 3 from e, there are two such edges; that means we have 4 choices of two vertices from X, each of which makes two quadrilaterals with e. That gives 120 quadrilaterals of this type. As for a negative quadrilateral with three negative edges, it is simply a path of length 3 in P, completed by a non-Petersen edge. There are 60 such paths. Thus we have a total of 180 negative quadrilaterals.

	$+P$	$-P$	P_D	$-P_D$
$+C_5$'s	12	0	6	6
$-C_5$'s	0	12	6	6
$+C_6$'s	10	10	0	0
$-C_6$'s	0	0	10	10

	$+K_{10}$	$-K_{10}$	K_P	$-K_P$
$+C_3$'s	120	0	60	60
$-C_3$'s	0	120	60	60
$+C_4$'s	630	630	450	450
$-C_4$'s	0	0	180	180

TABLE 3.1. Census of short circles in Petersen signed graphs and K_{10}.

I will define three more signatures on P in Section 8. In a sense they and $+P$, $-P$, and P_D are all that exist; see Theorem 8.1.

The circle compatibility graph.

Here are two important observations about the even circles in P, all of which are hexagons and octagons. Let's define two circles to be *compatible* if their intersection does not include a path of length 2. This relation defines a *compatibility graph* on $\mathcal{C}(P)$, the class of all circles in P, and in particular on the class $\mathcal{C}_l(P)$ of circles of length l.

Lemma 3.3. *Each vertex v determines a hexagon $H(v)$ which is the unique Hamiltonian circle in $P \setminus N[v]$. The mapping H is an isomorphism of P onto the compatibility graph of $\mathcal{C}_6(P)$.*

Proof. If vw is an edge, $H(v) \cap H(w)$ is $P_1 \cup\!\!\!\cdot\; P_1$. If v and w are not adjacent, $H(v) \cap H(w) = P_2$. □

Lemma 3.4. *Each edge vw determines an octagon $O(vw)$ which is the unique Hamiltonian circle in $P \setminus \{v, w\}$, and this gives a bijection $E(P) \leftrightarrow \mathcal{C}_8(P)$. No two octagons are compatible. The octagon $O(vw)$ and hexagon $H(x)$ are compatible if and only if x has distance 2 from v and w.*

Proof. The first part is true by inspection. Deleting two adjacent vertices leaves divalent vertices that force the Hamilton circle. Deleting two nonadjacent vertices creates a leaf so no octagon remains.

There are two possible relationships between two octagons $O(e)$ and $O(f)$. If e and f are adjacent, then $O(e) \cap O(f) = P_1 \cup\!\!\!\cdot\; P_1 \cup\!\!\!\cdot\; P_2$. The intersection if e and f are not adjacent is $P_1 \cup\!\!\!\cdot\; P_3$.

There are three possible relationships between an octagon $O(vw)$ and a hexagon $H(x)$. If $x = v$, $O(vw) \cap H(x)$ is $P_2 \cup\!\!\!\cdot\; P_2$. If $x \in N(v)$ (and $x \neq w$), $O(vw) \cap H(x) = P_4$. Otherwise, v has distance 2 from v and w; then $O(vw) \cap H(x) = P_1 \cup\!\!\!\cdot\; P_1$, which makes the circles compatible. □

The four vertices at distance 2 from vw are independent in P; therefore, by Lemma 3.3 they are incompatible. Thus, the graph of compatibilities amongst the hexagons and octagons is P with an extra vertex v_e adjoined for each edge e, with neighborhood the set of vertices at distance 2 from e. The degree of each P vertex in this graph is $3 + 4 = 7$ while the degree of v_e is 4. The graph has girth 4.

Part II. Balance and Imbalance

4. Switching

An essential operation on a signed graph is *switching*. There are two ways to express it. Switching a vertex set $X \subseteq V$ in Σ means reversing the signs of all edges between X and its complement X^c. In particular, switching a vertex v is the same as switching the set $\{v\}$, switching X is equivalent to switching all of its vertices, one after the other, and switching X^c is equivalent to switching X. Now, take a *switching function* $\tau : V \to \{+1, -1\}$. Switching Σ by τ means replacing the signature σ by a new signature σ^τ whose definition is $\sigma^\tau(e_{ij}) := \tau(v_i)^{-1}\sigma(e_{ij})\tau(v_j)$. (I inserted the inversion to suggest that switching is like conjugation—which it is.) Evidently, switching by τ is the same as switching by $-\tau$, and also as switching either $\tau^{-1}(-1)$ or $\tau^{-1}(+1)$.

A notation for switching that is often convenient is to list the switched vertex set. For $U \subseteq V$, define the signed characteristic function $\tau_U(v) := -$ if $v \in U$, $+$ if $v \notin U$. Then Σ^U is short for Σ^{τ_U}. That is, $\sigma^U(e) = \sigma(e)$ if the endpoints of e are both in U or both not in U, but if one endpoint is in U and the other is not, then $\sigma^U(e) = -\sigma(e)$.

The relation defined by $\Sigma \sim \Sigma'$ if Σ' is obtained by switching Σ is clearly an equivalence relation on signed graphs. I say Σ and Σ' are *switching equivalent*. An equivalence class is called a *switching class*.

Switching preserves the signs of circles and therefore preserves balance. The implication is two-way.

Proposition 4.1 (Sozański [24], Zaslavsky [29])**.** *Two signatures of the same underlying graph are switching equivalent if and only if they have the same list of positive circles.*

A coarser equivalence relation is $\Sigma \simeq \Sigma'$, meaning that Σ' is isomorphic to a switching of Σ. I call Σ and Σ' *switching isomorphic* and I call an equivalence class a *switching isomorphism class*. (I should say that most authors do not make this distinction between switching equivalence and switching isomorphism. The difference is only apparent when one has labelled graphs. Graph types, like P and K_P for instance, are unlabelled unless one specifies a vertex labelling, so there is no difference between equality and isomorphism.) It is plain to see that two signed graphs can only be switching isomorphic if their underlying graphs are isomorphic. The real meaning of switching isomorphism is in terms of the signs of circles.

Proposition 4.2. *If τ is a switching function on Σ and ψ is an isomorphism $|\Sigma| \to |\Sigma'|$, and if $(\Sigma^\tau)^\alpha = \Sigma'$, then α preserves the signs of circles. Conversely, if α is such an isomorphism $|\Sigma| \to |\Sigma'|$, there is a switching function τ such that $(\Sigma^\tau)^\alpha = \Sigma'$, which is unique up to negation on components of Σ.*

Theorem 4.3. *$+P$, $-P$, and P_D are not switching isomorphic. However, K_P and $-K_P$ are switching isomorphic, and P_D and $-P_D$ are switching isomorphic.*

The switching isomorphisms explain why the census of signed circles in Table 3.1 is the same for P_D and $-P_D$ as well as for K_P and $-K_P$.

Proof. The first part is elementary, my dear reader, as $+P$ is balanced but $-P$ is not. As for P_D, its census of signed pentagons or of hexagons (Table 3.1) shows it cannot switch to either $+P$ or $-P$. A direct proof is that it has a positive pentagon with vertices $12, 34, 15, 23, 45$, in which it differs from $-P$, and a negative hexagon with vertices $15, 24, 35, 14, 25, 34$, differing from $+P$.

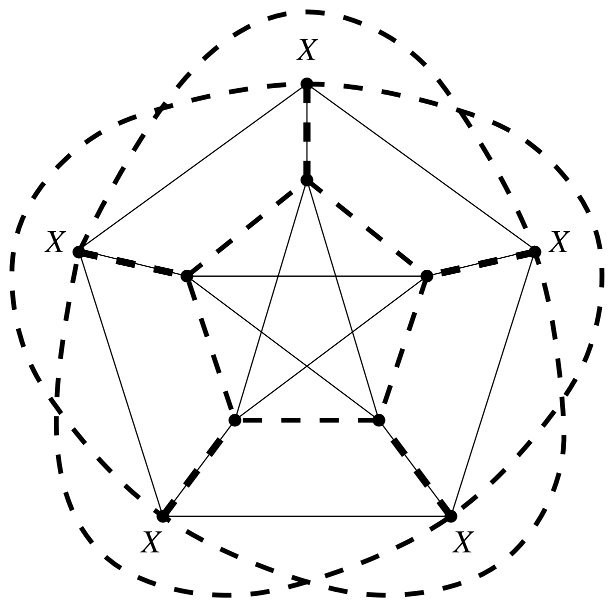

Figure 4.1. How to get K_P from $-K_P$ by switching. Solid lines: the positive edges of $-K_P$ before switching. Dashed lines: the negative edges of $(-K_P)^X$, i.e., after switching the five outer vertices; since they are P, $(-K_P)^X \simeq K_P$.

There is a cute way to switch $-K_P$ to K_P. P is made up of two pentagons C_5 (not at all uniquely, of course), joined by a perfect matching; this is obvious in the conventional drawing with one pentagon inside and the other one outside. Let $X :=$ the vertex set of one of those pentagons and switch $-K_P$ by X. Switching reverses the signs of all edges between the inner and outer pentagons. If we have a drawing of the positive subgraph of $-K_P$ (see Figure 4.1), the negative subgraph of $(-K_P)^X$ consists of the five matching edges and the complements of the inner and outer pentagons. But the complement of a pentagon is another pentagon. Thus the negative subgraph after switching is precisely P.

A similar argument applies to $-P_D$. Figure 4.2 shows the vertices to switch in $-P_D$, and the vertex relabelling, to get P_D. □

Having shown that K_P switches to be isomorphic to its own negation, as does P_D, we might like to know whether a labelled K_P is switching equivalent to $-K_P$, and similarly for P_D and $-P_D$. This is a more stringent question: we are asking whether the signs of all edges of K_P, or P_D, can be negated by switching. They cannot. This illustrates a general fact.

Proposition 4.4. *A signed graph switches to its negation if and only if it is bipartite.*

Proof. Supposing Σ is bipartite, switching one of the two color classes negates every edge.

If we want to negate every edge, then the switching set X has to contain exactly one endpoint of each edge. There is such an X if and only if the graph is bipartite. □

Note that this switching does not involve an isomorphism.

We've seen how to represent a signed K_n by its positive or negative subgraph. This was how Seidel thought of switching: as changing the adjacencies in a graph by reversing

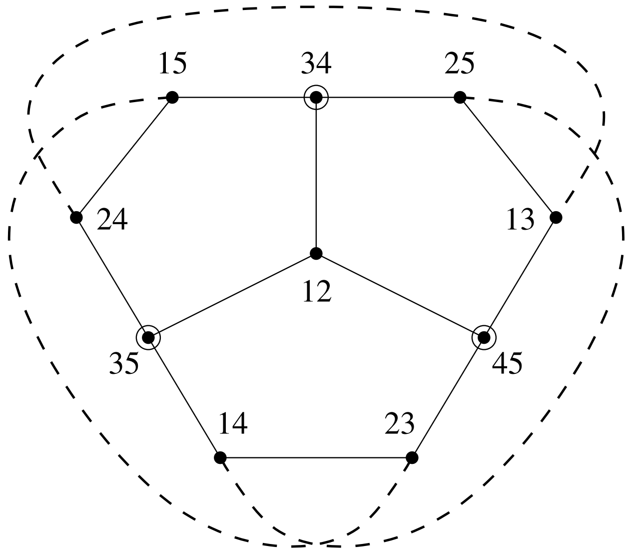

FIGURE 4.2. How to get $-P_D$ from P_D: Switch the neighbors (circled) of the vertex 12 farthest from the special hexagon.

those between a vertex set X and X^c; this operation is now called *Seidel switching* or *graph switching*. Some of the graphs that are obtained from P by graph switching, or from the signed standpoint, which appear as the negative subgraphs of switchings of K_P, are shown in Figure 4.3; another is $-K_P$, as we saw in Theorem 4.3.

The graph of Figure 4.3 in which a vertex is isolated—i.e., its edges are all positive in the signed complete graph—call it $K_1 \mathbin{\dot\cup} \Gamma_9$, is especially interesting. (By symmetry of P, there is only one such graph, up to isomorphism.) Γ_9 is strongly regular; it is the Paley graph of order 9 [12, Section 10.3]. I will discuss this more deeply in Section **??**.

Finally, in Figure 4.3 I show some of the graphs obtained by graph switching of P, i.e., the negative subgraphs of some switchings of K_P. Notice that there are two ways to get K_P back: X or X^c must be either $\varnothing$ or $N(u, w)$ for vertices that are adjacent in P. This will be proved in Lemma 6.5.

5. Balance and Antibalance

The fundamental concept of signed graph theory is the sign of a circle C. This is the product of the signs of its edges, written $\sigma(C)$. Switching does not change the sign of a circle; in fact, it is easy to prove that two different signatures on the same graph have the same circle signs if and only if they are related by switching. Thus, a fundamental property of a signed graph is whether its circles all have positive sign. If they do, we call it *balanced*. The obvious balanced graph is $+\Gamma$, with all edges positive; but any switching of $+\Gamma$ is also balanced. The first theorem of signed graphs, due to Frank Harary, involves switching implicitly:

Theorem 5.1 (Harary [13]). *Σ is balanced if and only if V can be bipartitioned into subsets V_1 and V_2 so that every negative edge has one endpoint in each subset and each positive edge has both endpoints in the same subset.*

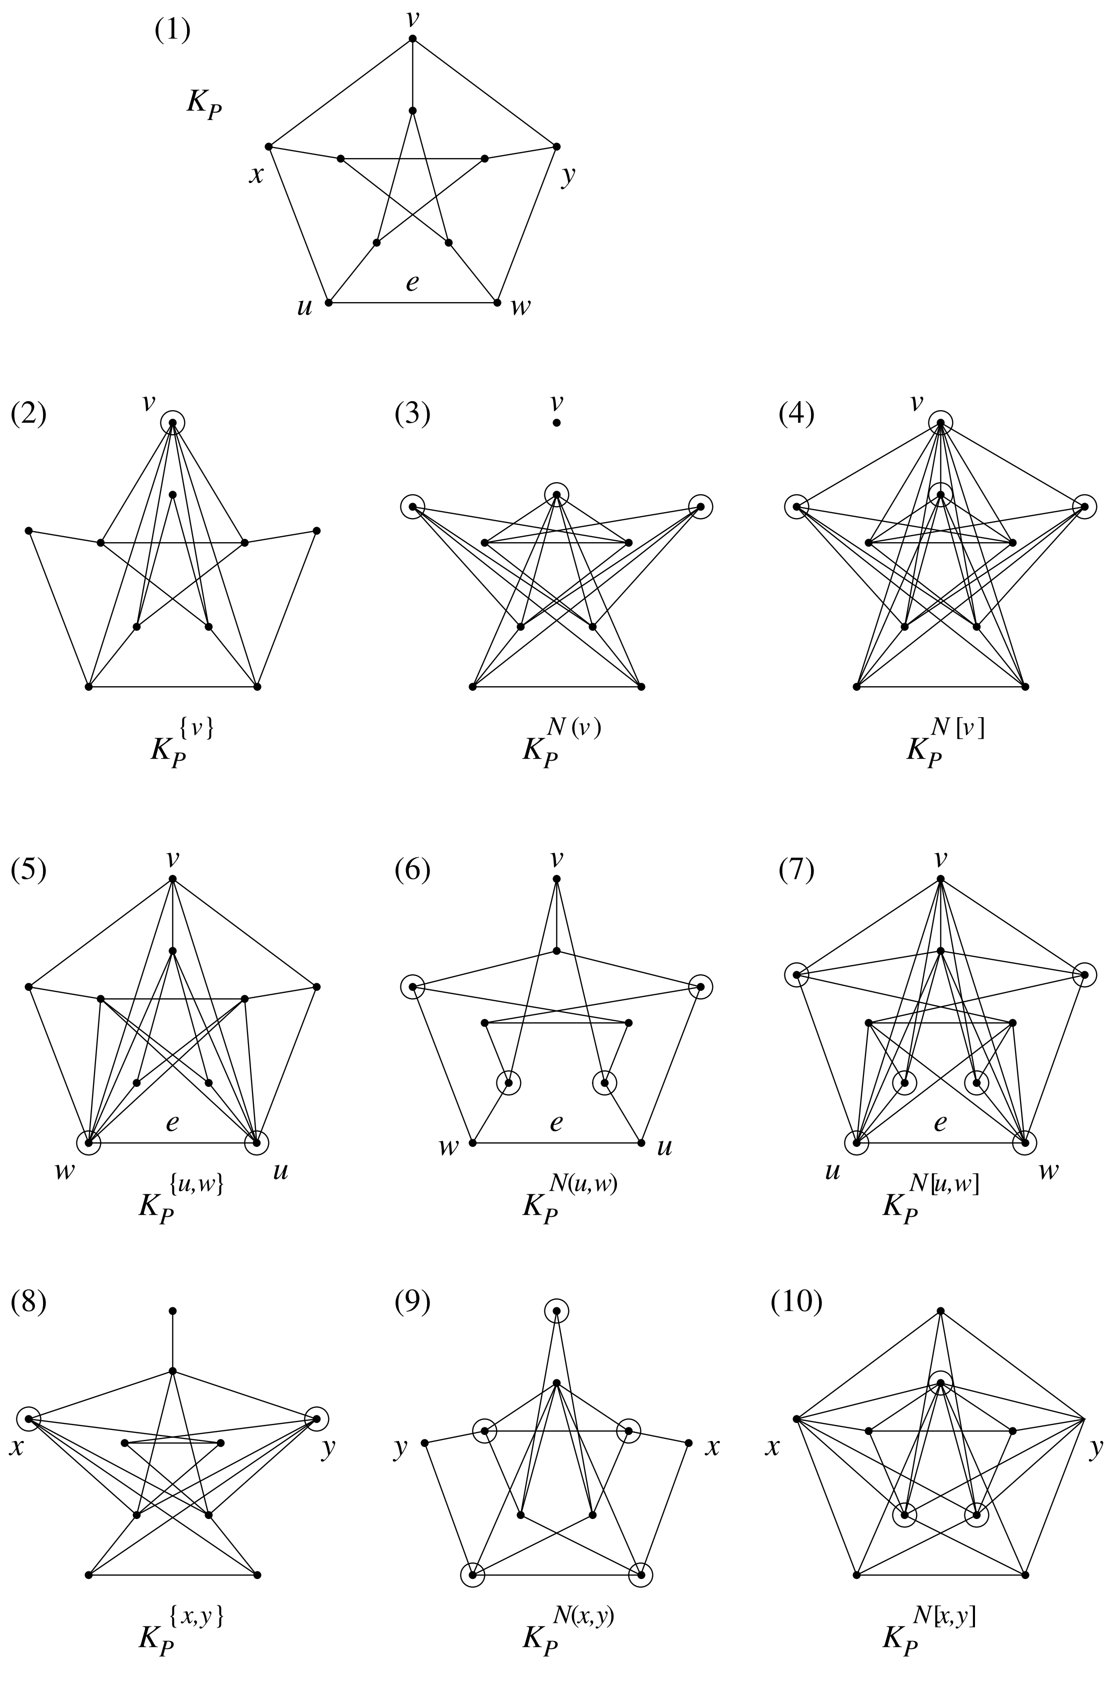

FIGURE 4.3. Some graphs that are negative subgraphs of switchings of K_P. Switched vertices are circled, except in $K_P^{N[x,y]}$ where the unswitched vertices are circled.

I call such a bipartition a *Harary bipartition* of Σ. Stated in terms of switching: Σ is balanced if and only if it can be switched to be all positive. (This is implicit in work of E. Sampathkumar [20] and explicit in [30, Cor. 3.3].)

A further characterization is:

Theorem 5.2 (Harary [14]). *A signed 2-connected graph is balanced if and only if every circle through a fixed vertex is positive.*

The Petersen signed graphs, $+P$, $-P$, K_P, and $-K_P$, are all 2-connected, and only the first is balanced. We can conclude either from Theorem 5.2 or the symmetry of P that in each of the latter three, every vertex lies on some negative circle.

An opposite property of Σ is balance of $-\Sigma$. That is, if we negate all the signs in Σ, do we get a balanced graph? If we do, we call Σ *antibalanced.* More directly, Σ is antibalanced when circles of even length are positive and those of odd length are negative; in a formula, $\sigma(C) = (-1)^{|C|}$. The canonical example is $-\Gamma$, where all edges are negative. It is fair to say that antibalance is the signed-graph generalization of biparticity. From Theorem 5.1 it is obvious that:

Corollary 5.3. *Σ is antibalanced if and only if $V = V_1 \cup V_2$ in such a way that every negative edge crosses between V_1 and V_2 and each positive edge has both endpoints in one of V_1 and V_2.*

Corollary 5.4. *For a signed graph, any two of the following properties imply the third: balance, antibalance, and biparticity of the underlying graph.*

6. Measurement of Imbalance: Frustration

As balance is the fundamental concept of signed graph theory, we want ways to measure how far from balance a signed graph is. Probably the commonest and most useful measure is the *frustration index* $l(\Sigma)$, which is defined as the smallest number of edges whose deletion leaves a balanced subgraph—or, equivalently, whose sign reversal makes the graph balanced (the equivalence was proved by Harary (1959b)). An edge set whose deletion leaves a balanced subgraph is a *balancing set* of Σ. A *balancing edge* is a single edge that constitutes a balancing set in an unbalanced signed graph. Determining frustration index includes the Max-Cut problem, i.e., finding the largest edge set of the form $E(X, X^c)$ in a graph, since a maximum cut in Γ is the complement of a minimum balancing set in $-\Gamma$ (see Lemma 6.1); it follows that frustration index is NP-hard. However, it is clear that $l(\Sigma) \le |E^-|$, and there are stronger results.

Lemma 6.1. *The minimal balancing sets in $-\Gamma$ are the complements of the maximal cuts in Γ.*

Proof. A bipartite subgraph in $-\Gamma$ is precisely what is a balanced subgraph. A minimal balancing set is the complement of a maximal bipartite subgraph. A maximal bipartite subgraph is a maximal cut. □

Lemma 6.2. $l(\Sigma) = \min_X |E^-(\Sigma^X)|$.

Proof. Let S be a minimal balancing set. (It need not have minimum cardinality.) Since $\Sigma \setminus S$ is balanced, it switches to $\Sigma^X \setminus S$ which is all positive. Restore S with the same switching. Every edge of S must be negative, or else there would be a smaller balancing set. □

Proposition 6.3. *If every cut in Σ has at least as many positive as negative edges, then $l(\Sigma) = |E^-|$. If some cut has more negative than positive edges, then $l(\Sigma) < |E^-|$.*

Proof. If $|E^-(X, X^c)| > |E^+(X, X^c)|$, then switching X reduces the number of negative edges. If $|E^-(X, X^c)| \leq |E^+(X, X^c)|$ for every X, then no switching can reduce the number of negative edges; so $l(\Sigma) = |E^-|$ by Lemma 6.2. □

Some Petersens.

Let's apply Proposition 6.3 to Petersen signed graphs. Of course, for any graph $l(+\Gamma) = 0$, so I'll omit the all-positive Petersen, and $l(P(e)) = 1$. A structural lemma for $-P$ is helpful.

Lemma 6.4. *$-P$ switches to have just three negative edges, all at mutual distances 3.*

Proof. Switch an independent set of four vertices (which is the four neighbors of an edge) and observe the result. □

Lemma 6.5. *A switching K_P^X is isomorphic to K_P if and only if X or X^c is either $\varnothing$ or $N(u, w)$ for some edge uw of the Petersen graph. For every other X, $|E^-(K_P^X)| > 15$.*

Proof. Sufficiency is visible in Figure 4.3.

Now, consider the effect of switching X on the number of negative edges. Here we apply the method of Proposition 6.3. We have to examine all the cuts $E(X, X^c)$, where $E := E(K_{10})$, and show that $|E(P)(X, X^c)| \leq \frac{1}{2}E(X, X^c)$ in every case. If $|X| = k \leq 3$, $|E(X, X^c)| = k(10 - k)$ but $|E(P)(X, X^c)| \leq 3k < \frac{1}{2}k(10 - k)$. If $k = 5$, $|E(X, X^c)| = 25$, while $|E(P)(X, X^c)| \leq 15 - 4 < \frac{1}{2}(25)$ because the subgraph of P induced by X contains at least 2 edges. (We know this from Lemma 3.2. A maximum independent vertex set has 4 vertices and has the form $V(i) := \{ij : j \neq i\}$. Adding a fifth vertex creates at least 2 edges within X. If X has fewer than 4 independent vertices, it also contains at least 2 edges.) That leaves $|X| = 4$. Then $|E(X, X^c)| = 24$ and $|E(P)(X, X^c)| \leq 3 \cdot 4 = \frac{1}{2}24$. Thus, K_P^X has at least 15 negative edges, and the only way it can have as few as 15 is for $|E(P)(X, X^c)| = 3 \cdot 4$, so X has to be independent. □

Proposition 6.6. *$l(-P) = 3$ and $l(K_P) = l(-K_P) = |E(P)| = 15$, and the minimum balancing sets are unique up to symmetry of P.*

Proof. Each graph has to be treated separately, except that $l(-K_P) = l(K_P)$ because they are switching isomorphic, and $l(-P_D) = l(P_D)$ similarly. The good part is that $-P$ and K_P illustrate different approaches.

The frustration index of $-P$ follows from the known fact that the maximum size of a cut in P is 12; but it will be more interesting to show directly why $l(-P) = 3$ and to find the structure of all minimum balancing sets.

In $-P$ there are four pentagons on each edge, all of which are negative. There are 15 edges and 5 edges in each pentagon; it follows that there are 12 pentagons. By deleting k edges we can eliminate at most $4k$ pentagons. Therefore, no balancing set can be smaller than 3 edges. A balancing set of three edges e, f, g must not have two of its edges in one pentagon. If we look closely, we see that the four pentagons on e cover all but two edges of P. These two edges are the only possible choices for f and g. Note that these edges all have distance 3 from each other; we have seen that this configuration is unique up to symmetries of P (Lemma 3.1). Let's delete all three edges. We get a bipartite graph, which is balanced (by Corollary 5.4) because it is all negative. Figure 6.1 shows $-P$ switched so that, as promised by Lemma 6.2, only the three edges of a minimum balancing set are negative.

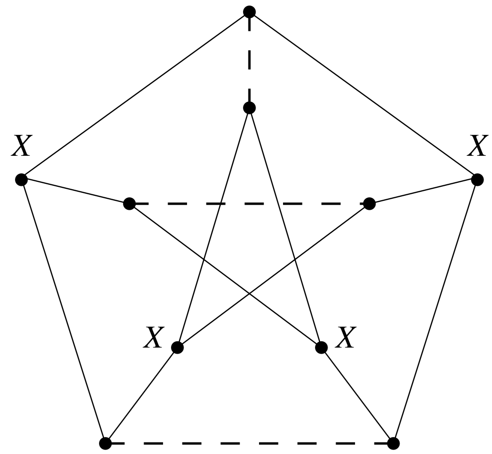

FIGURE 6.1. $-P$, switched by X so that E^- is a minimum balancing set (of three edges). Notice that X is a maximum independent vertex set.

The minimum balancing set B is unique once the first edge is chosen. Since P is edge transitive, B is unique up to symmetry.

The other example, K_P, is solved by Lemma 6.5. □

The switching behavior of P_D is unexpectedly complex. In a hexagon $H = uvwxyzu$ label the edges alternately a and b. There are two kinds of pentagon: those whose intersection with H contains one a edge and either 0 or 2 b edges, and those with the opposite property. Let $\mathcal{C}_5^a(H)$ be the set of pentagons of the first type.

Lemma 6.7. *The bipartition* $\{\mathcal{C}_5^a(H), \mathcal{C}_5^b(H)\}$ *is independent of the hexagon* H*.*

According to Lemma 6.7 we can unambiguously label every hexagon with an a triple and a b triple so that $\mathcal{C}_5^a(H)$ is the same for all hexagons. Therefore we can drop the H in the notation and simply divide the pentagons into the class $\mathcal{C}_5^a$ of a pentagons and the class $\mathcal{C}_5^b$ of b pentagons. Which class is a and which is b is arbitrary. In any particular exemplar of P_D, the three negative edges are either an a triple or a b triple of P; let us call these two types of signed graph P_D^a and P_D^b. Write $P_D^a(H)$ for the particular signature in which the a edges of the hexagon H are the negative edges, and similarly for $P_D^b(H)$.

Lemma 6.8. *All* P_D^a *signatures are switching equivalent, as are all* P_D^b *signatures, but* $P_D^a \not\sim P_D^b$*. Furthermore,* $-P_D^a \sim P_D^b$ *and* $-P_D^b \sim P_D^a$*.*

That is, by switching P_D^a one can make any a triple the negative edge set, but no other triple and in particular no b triple.

Proof of Lemmas 6.7 and 6.8. Start with a fixed hexagon H_0 and the signature $P_D^a(H_0)$. Consider a hexagon H, adjacent to H_0 in the compatibility graph of P (Section 3). The hexagons intersect in two opposite edges of each, one of which is negative. Let that be e, and switch $V(e)$. The result is $P_D^a(H)$, as in Figure 6.2(a). This shows that by switching we can transform $P_D^a(H_0)$ into $P_D^a(H)$ for any adjacent pair of hexagons, and by repeated switching for any pair of hexagons since the hexagon compatibility graph is connected.

Now, consider what distinguishes the a edges in H_0. They are the edges f such that the pentagon P_0 that meets H_0 in only the one edge f, or a three-edge path of which the middle edge is f, is negative. It follows that $\mathcal{C}_5^a$ is the class of negative pentagons, and $\mathcal{C}_5^b$ is the class of positive pentagons. This is invariant under switching, so only the a triples can be made into negative triples by switching. We have proved Lemma 6.7.

Finally, consider the switching of $-P_D^a$ shown in Figure 6.2(b). This proves $-P_D^a(H)$ switches to $P_D^b(H)$. □

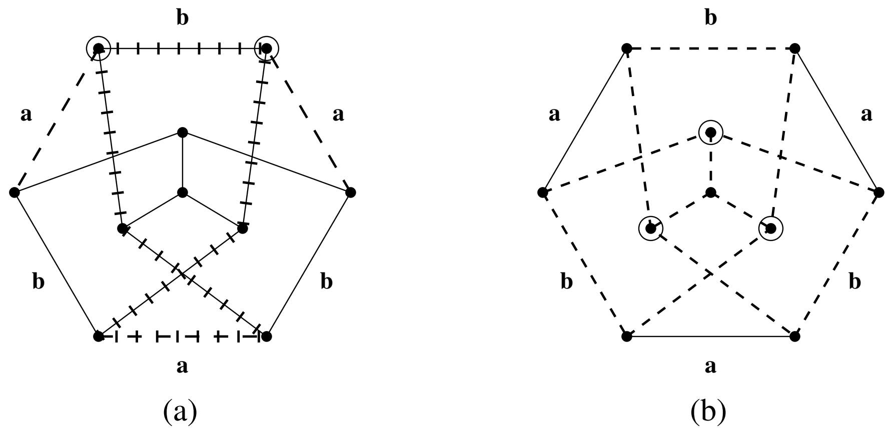

FIGURE 6.2. The vertices (circled) to switch to change (a) $P_D^a(H_0)$ into $P_D^a(H)$ where H_0 (the outer hexagon) and H (railroad tracks) are adjacent, and (b) $-P_D^a(H_0)$ into $P_D^b(H_0)$. The a and b edges of H_0 are marked.

Theorem 6.9. $l(P_D) = l(-P_D) = 3$. *There are* 12 *minimum balancing sets; they are the* a *triples (if* P_D^a*), or else the* b *triples (if* P_D^b*), in the twelve hexagons.*

Proof. As with K_P, $l(-P_D) = l(P_D)$ because $-P_D \simeq P_D$.

The upper bound of 3 on the frustration index of P_D is visible in Figure 3.2. The index cannot be < 2 because an edge is in exactly four pentagons, so no one negative edge can give the six negative pentagons of P_D. Suppose there is a minimum balancing set S of two edges at distance 2. They cannot be adjacent, or else switching the common vertex would give a smaller balancing set. If their distance is 2 they generate six negative hexagons, and if it is 3 they generate four negative hexagons; therefore they cannot switch to P_D. Thus, $l(P_D) > 2$, which solves $l(P_D)$.

The rest follows from Lemmas 6.7 and 6.8, once we prove there are no other balancing sets of three edges than the a or b triples. To settle this, consider the effect on P_D of switching a set X so as to keep the number of negative edges at 3. Let H_0 be the special hexagon (see Figure 6.3). Call vertices *equivalent* if they are both in X or both in X^c.

The cut $E(X, X^c)$ must have equally many positive and negative edges, and up to 3 negative edges. Negating one negative edge will increase $|E^-|$ because P is edge 3-connected.

If switching negates exactly two negative edges, those edges must have exactly one endpoint in X and the third must have either no endpoints or both in X. For definiteness, say the switched negative edges are uv and yz, so their endpoints are inequivalent, while those of wx are equivalent. The picture will look like Figure 6.3(a) or (b, c), where X vertices are circled and X^c vertices are boxed. The difference is that in (a) the endpoints of uz are inequivalent, while in (b, c) they are equivalent. In (b) the endpoints of vw are equivalent,

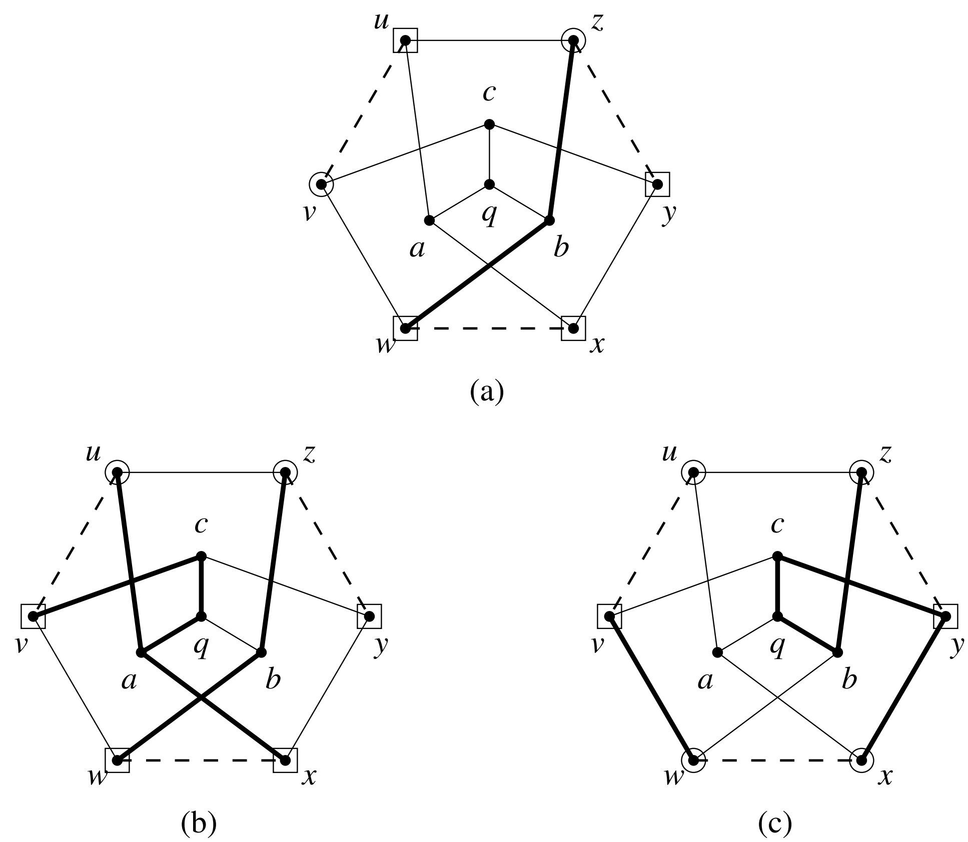

FIGURE 6.3. Switching P_D to maintain 3 negative edges of which wx is originally negative.

in (c) they are inequivalent. The edges of H_0 that switch to become negative, if any, are drawn with heavy lines.

To test these configurations we look at paths in $P_D^+ \setminus H_0$ between inequivalent vertices. Any such path contains at least one edge that becomes negative after switching. Thus, the number of such paths is a lower bound for the number of new negative edges. If this bound is too high, after switching $|E^-| > 3$, and the configuration is ruled out. In (a) we may without loss of generality assign w and x to X^c. Then the path wbz makes for a fourth negative edge after switching, along with the unswitched edge wx and the switched edges vw and uz. This case, therefore, is impossible. In (b) the three heavy paths $uaqcv$, uax, and zbw connect inequivalent vertices, so they each contain a new negative edge. Since wx remains negative, there are only two negative edges for the three paths; thus, ua must be switched. Similarly, zb must be switched. This tells us that $X = \{u, z\}$, or else there will be more than three negative edges after switching. Case (c) is impossible because after switching there are already three negative edges in H_0 and one more in the heavy path $ycqbz$.

Suppose all three negative edges in H_0 are made positive by switching. Then X must occupy alternating vertices around H_0, as in Figure 6.4(a), or not, in which case we have the configuration of Figure 6.4(b). The first of these is ruled out by the need for a fourth negative edge in the path vcy, after switching. The second is a possible case. H_0 switches to have one negative edge, vw. The five paths uax and $wbqcv$, $wbqcy$, $zbqcv$, $zbqcy$ all need

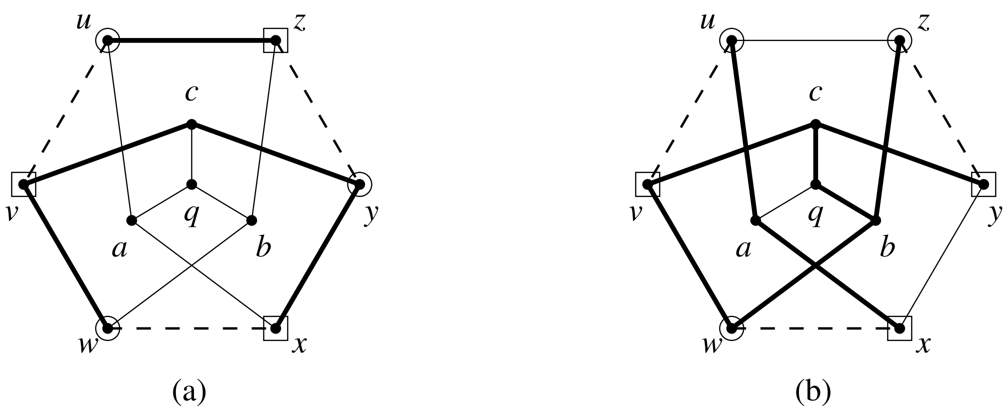

Figure 6.4. Switching P_D to keep 3 negative edges.

to have a negative edge while using only two negative edges among them; thus, ua or xa must be negative, and qb or qc must be negative, after switching. The path $uaqcy$ shows that these edges cannot be xa and qb. The path $xaqbz$ shows they cannot be ua and qc. Thus, either ua, qb are negative, or xa, qc are negative. These two cases are symmetrical by the reflection $u, v, w \leftrightarrow z, y, x$, so we take only the first of them.

Here $X = \{u, z, b, w\}$. This is the vertex set of a path in P_D^+ of length 3, starting with a positive edge in H_0 and ending at an opposite vertex of H_0. Now, in $P_D^{\{u,z\}}$, the special hexagon has bw as a positive edge, so it can be switched to produce an isomorph of P_D. Thus, viewing P_D^X as $(P_D^{\{u,z\}})^{\{b,w\}}$. Thus, we have a proof of Theorem 6.9 as well as a complete description of the switchings that produce copies of P_D. □

Proposition 6.10. *$P_D^X \cong P_D$ if and only if $X = \varnothing$, or X is the vertex set of a positive edge in the special hexagon of P_D, or X is the vertex set of an all-positive path from a positive edge in the special hexagon to another vertex in the special hexagon.*

The process of switching the vertex set of a positive edge in the special hexagon will transform any P_D^a to any other in at most two steps. □

7. Maximum Frustration

We found two unbalanced signatures of P whose frustration index is 3. Could this be the maximum frustration possible for any signature? The maximum frustration of signatures of a simple graph has been studied at length by Akiyama, Avis, Chvátal, and Era (1981a) and subsequent writers, but it is hard; almost all results are only asymptotic, which will not help us with our Petersen examples.

The most elementary general statements are upper bounds. Write $l_{\max}(\Gamma) := \max_\sigma l(\Gamma, \sigma)$ for the maximum frustration index of signatures of Γ.

Proposition 7.1. *A minimum balancing set S of Σ satisfies $\deg_S(v) \leq \frac{1}{2}\deg_{|\Sigma|}(v)$ at every vertex. Thus, $l_{\max}(\Gamma) \leq \frac{1}{2}|E|$.*

Proof. Assume by switching that $S = E^-$. If $\deg_S(v) > \frac{1}{2}\deg_{|\Sigma|}(v)$ at a vertex, switch v. This reduces the size of E^-, but that gives a smaller balancing set. Thus the first statement is true. The second is an immediate consequence. □

This is a weak result, even though useful. A vertex star is a cut; a similar conclusion holds for every cut.

For a cubic graph like P Proposition 7.1 means that $l_{\max}(\Gamma) \leq \frac{1}{2}|V|$ and a minimum balancing set is a matching (not necessarily a perfect matching). For a signed K_{10} it means $l_{\max}(K_{10}) \leq 20$.

Almost the only known general exact solution is Petersdorf's.

Proposition 7.2 (Petersdorf (1966a))**.** *For complete graphs there is the exact upper bound*

$$l_{\max}(K_n) = \left\lfloor \left(\frac{n-1}{2}\right)^2 \right\rfloor.$$

Furthermore, $l(K_n, \sigma)$ achieves this value if and only if (K_n, σ) is antibalanced, e.g., $-K_n$.

Proof. First, $-K_n$. A maximum cut in K_n is $e(X, X^c)$ where $|X| = \lfloor (n+1)/2 \rfloor$. The complement of $E(X, X^c)$ has $\lfloor (\frac{1}{2}(n-1))^2 \rfloor$ edges, which is therefore the value of $l(-K_n)$.

Now we have to prove this is the maximum. Suppose $\Sigma := (K_n, \sigma)$ is signed so that $E^- = l_{\max}(K_n) = \max_\sigma l(K_n, \sigma)$. By Proposition 7.1, $\deg_{\Sigma^-}(v) \leq \lfloor (n-1)/2 \rfloor$ for every vertex. That means $|E^-| \leq \frac{1}{2} n \lfloor (n-1)/2 \rfloor$.

If n is even, $|E^-| \leq n(n-2)/4 = \lfloor (\frac{1}{2}[n-1])^2 \rfloor$.

The odd case $n = 2k+1$ is more complicated. Suppose no two vertices of negative degree k are adjacent. Then at most $k+1$ vertices have degree k and at least k have degree $k-1$ in Σ^-. Therefore, $|E^-| \leq \frac{1}{2}[(k+1)k + k(k-1)] = ([n-1]/2)^2$. In the opposite case there are non-neighboring vertices v, w of degree k. By switching v we change k negative edges at v to positive and k positive edges to negative, thereby not changing $|E^-|$. However, we do increase the degree of w to $k+1 > \frac{1}{2}n$, which means that by switching w we can reduce the size of $|E^-|$. By Proposition 7.1, this contradicts the hypothesis that $|E^-(\Sigma)| = l(K_n, \sigma) = \max_\sigma l(K_n, \sigma)$. Therefore, no two such vertices can exist, which proves that $-K_n$ does indeed maximize $l(K_n, \sigma)$.

I omit proving that the switching class of $-K_n$ uniquely maximizes frustration index. The proof so far contains hints of how to proceed by comparing degrees in the negative subgraph. □

For K_{10} the exact bound is indeed 20 as Proposition 7.1 suggests. K_P falls well short of that.

Let's examine the Petersen graph.

Proposition 7.3. $l_{\max}(P) = 3$. *Any signed Petersen graph with frustration index* 3 *is switching isomorphic to $-P$ or P_D.*

Proof. Assume we have (P, σ) with maximum frustration index, switched so that E^- is a minimum balancing set. Thus, E^- is a matching of some size ≤ 5 (and at least 3).

We dispose of $|E^-| = 5$ quickly. A 5-edge matching M is a cut separating two pentagons. (The proof is that $P \setminus M$ is 2-regular. Since P has girth 5 and is not Hamiltonian, $P \setminus M$ is the union of two pentagons.) If $|S| = 5$, switching one side of the cut makes the graph all positive; that is, $(, \sigma)$ is balanced. Thus, E^- cannot have 5 edges.

For smaller E^- we study the structure. Recall that we have (P, σ) with $|E^-| = l(P, \sigma)$.

Let's assume first that E^- contains a pair of edges at distance 2. Up to symmetry there is only one way to have them, since they are the end edges of a 4-arc and P is 4-arc transitive. (See Figure 7.1(a).) To fix the notation, let them be uv, wx where v, w are adjacent and the other neighbors of v and w are u, u' and x, x', respectively. Any edge incident to any of u, v, w, x cannot be in E^-. The same holds for u', x', for the following reason. Switching $\{v, w\}$ replaces uv, wx in E^- by $u'v, wx'$ without changing its size; thus the switched negative edge set cannot be smaller than $|E^-|$. If E^- contained an edge $u't$, say, then the switched negative edge set would contain two adjacent edges $u't, u'v$ and therefore be reducible by switching, contradicting the assumption that $|E^-| = l(P, \sigma)$. That leaves only two possible edges for E^-: a, b in Figure 7.1(a). But if $a \in E^-$, then switching $\{u, x, y\}$ reduces the number of negative edges, contrary to assumption. If $b \in E^-$, the negative edges are alternate edges of a hexagon so we have P_D.

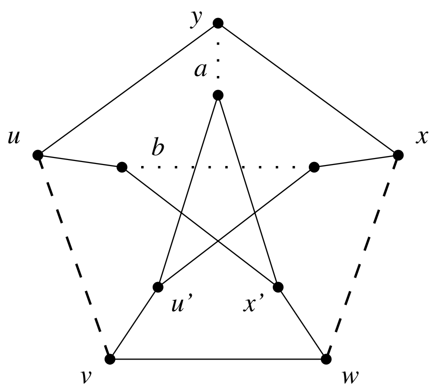

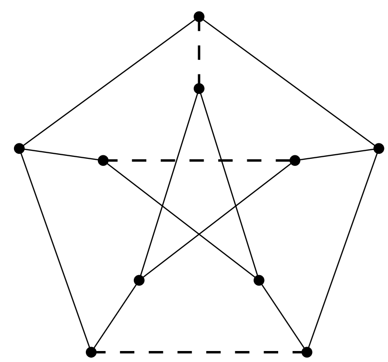

FIGURE 7.1. Possible edge pairs in a largest minimal balancing set of a signed P. (a) Negative edges, uv and wx, at distance 2. (b) No such edges.

If E^- has no two edges as close as distance 2, by Lemma 3.1 it can only be a set of three edges with mutual distance 3, as in Figure 7.1(b). This switches to $-P$ (see Figure 3.2).

We have proved that a signed P has maximum frustration 3 and the only examples are P_D and $-P$ (up to switching). $\square$

8. A CENSUS OF SIGNED PETERSEN GRAPHS

With the tools we have now we can prove there are exactly six inequivalent ways, up to switching isomorphism, to sign P. We know three of them; the remaining three are

$P(e)$, with one negative edge;
$P_{2,2}$, with two negative edges at distance 2; and
$P_{2,3}$, with two negative edges at distance 3.

Theorem 8.1. *There are exactly six switching isomorphism types of signed Petersen graphs. They are $+P$, $-P$, $P(e)$, $P_{2,2}$, $P_{2,3} \simeq -P(e)$, and P_D. The frustration indices are as in Table 8.1.*

Proof. The frustration index of $P(e)$ is obvious, since $P(e)$ is unbalanced. The frustration index of $P_{2,2}$ or $P_{2,3}$ is not less than 2 because in each of them there are two vertex-disjoint negative pentagons.

The census of small circles in Table 8.1, which can be verified by inspection, shows that the six signed Petersens are switching nonisomorphic. It remains to show they represent all the switching isomorphism classes. That is obvious for frustration index 0 or 1. For index 3 or more it is Proposition 7.3.

Thus, let's assume we have (P,σ) with $l(P,\sigma) = 2$. We can assume by switching that there are two negative edges. If these edges are adjacent at v, switching v reduces $|D^-|$ to 1, contrary to our assumption on frustration index. If these edges are at distance 2, we have $P_{2,2}$. There is only one way to have edges at distance 2, because with the connecting edge they form a path of length 3, and P is 4-arc transitive. That shows $P_{2,2}$ is unique. Similarly, any two negative edges at distance 3 are symmetric, by Lemma 3.1. Thus $P_{2,3}$ is unique. □

$(P,\sigma) =$ $\simeq$	$+P$	$P(e)$ $-P_{2,3}$	$P_{2,2}$ $-P_{2,2}$	$P_{2,3}$ $-P(e)$	P_D $-P_D$	$-P$
$l(P,\sigma)$	0	1	2	2	3	3
$+C_5$'s	12	4	6	8	6	0
$-C_5$'s	0	8	6	4	6	12
$+C_6$'s	10	6	4	6	0	10
$-C_6$'s	0	4	6	4	10	0

TABLE 8.1. The six switching isomorphism classes of signed Petersen graphs (P,σ).

A kind of signature that one may naturally be curious about is $P(C_l)$, where the negative edges are a circle of length l, for $l = 5, 6, 8, 9$. Of course, they are switching isomorphic to signatures in Table 8.1.

Proposition 8.2. *$P(C_5) \simeq P_D$, $P(C_6) \simeq -P$, $P(C_8) \simeq +P$, and $P(C_9) \simeq P_{2,2}$.*

The proofs are a pleasant and easy exercise. For instance, $P(C_9) \sim -P_{2,2} \simeq P_{2,2}$.

9. MEASUREMENT OF IMBALANCE: VERTEX DELETION

The vertex version of frustration index is the *vertex deletion number* $l_0(\Sigma)$, the least number of vertices whose deletion leaves a balanced subgraph, or if you prefer, the complement of the largest order of a balanced induced subgraph. Call $X \subseteq V$ a *balancing vertex set* if $\Sigma \setminus X$ is balanced. There is much less literature on balancing vertex sets and the vertex deletion number than on edge frustration, possibly because the removal of an object is inappropriate to social and physical applications. Still, there is enough to say. The first and rather obvious facts are:

Proposition 9.1. *The vertex deletion number depends only on the switching class of Σ. It satisfies $l_0(\Sigma) \leq l(\Sigma)$.* □

It is apparent that $l_0(\Sigma) = 0$ if Σ is balanced and 1 if Σ has a balancing edge. To get a complete list of values we can use a result similar to Lemma 6.2. The *vertex cover number* of a graph is $\beta(\Gamma) :=$ the least number of vertices such that are collectively incident to every edge. The largest order of an induced bipartite subgraph is $b_0(\Gamma)$.

Lemma 9.2. *The vertex frustration number satisfies*

$$l_0(\Sigma) = \min_{U \subseteq V} \beta((\Sigma^U)^-).$$

In particular, $l_0(+\Gamma) = 0$ *and* $l_0(-\Gamma) = n - b_0(\Gamma)$.

Proof. We show that, for any vertex subset X, the deleted subgraph $\Sigma \setminus X$ is balanced if and only if X is a vertex cover of the negative subgraph $(\Sigma^U)^-$ in some switching Σ^U.

Sufficiency of the vertex cover property is obvious: $\Sigma^U \setminus X$ is all positive, hence balanced.

Conversely, suppose $\Sigma \setminus X$ is balanced. Choose $U \subseteq X^c$ so that $(\Sigma \setminus X)^U$ is all positive. Since $\Sigma^U \setminus X$ is all positive, X covers all negative edges of Σ^U.

The value for an all-negative graph depends on the fact that such a graph is balanced if and only if it is bipartite. □

Lemma 9.2 can be greatly strengthened for a signed complete graph. Let's switch K_Γ so that a vertex u becomes isolated in the negative subgraph. In order to do this, we switch exactly the neighbors of u in Γ.

Theorem 9.3. *The vertex deletion number of a signed complete graph* (K_n, σ) *satisfies*

$$l_0(K_\Gamma) = \min_{u \in V} \beta\big((K_\Gamma^{N_\Gamma(u)})^-\big).$$

Proof. We prove a slightly stronger fact: The minimal balancing vertex sets of K_Γ are precisely the vertex covers in switchings of K_Γ such that the negative subgraph has an isolated vertex.

A minimal balancing vertex set X is smaller than V, so there is a vertex $u \notin X$. Switch by $N_\Gamma(u)$, so u becomes isolated in the negative subgraph. Now X must cover every negative edge xy, since otherwise $\{u, x, y\}$ supports a negative triangle. On the other hand, for any vertex set Y that covers all negative edges of the switched graph, $K_\Gamma^{N_\Gamma(u)} \setminus Y$ is all positive, hence balanced. □

Proposition 9.4. *The vertex deletion numbers of Petersen signed graphs are as in Table 9.1.*

Σ	$+P$	$P(e)$	$P_{2,2}$	$P_{2,3}$	P_D	$-P$	K_P
$l_0(\Sigma)$	0	1	2	2	3	3	6

TABLE 9.1. The vertex deletion numbers of Petersen signed graphs.

Proof. The first two numbers are obvious.

The values for $P_{2,3}$ and $P_{2,2}$ are clear: $l_0 \leq l = 2$, and the two vertex-disjoint negative circles imply $l_0 \geq 2$.

Similarly, $2 \leq l_0(P_D) \leq 3$. Suppose $l_0(P_D) = 2$. Then there are vertices u, v such that, in some switching of P_D, u and v cover all the negative edges. By switching u and v as necessary we can ensure that neither has negative degree > 1. The inevitable conclusion is that $l(P_D) \leq 2$; but we know this is not so (Proposition 6.6).

The value for $-P$ derives from the fact that there are no two vertices whose deletion from P gives a bipartite graph. On the other hand, deleting the neighborhood of any vertex leaves $C_6 \cupdot K_1$, which is very bipartite.

To verify the vertex deletion number of K_P we examine $K_P^{N_P(u)}$ for every vertex. Of course, there is only one vertex up to symmetry, so this is easy. The resulting negative subgraph, shown in Figure 4.3(??), has $\beta = 6$. That solves K_P. (It's interesting that P itself has the same vertex cover number, so the right answer could be obtained without switching, but we could not prove it was correct.) $\square$

10. Measurement of Imbalance: Clusterability

Clusterability is a generalization of balance based on Harary's Balance Theorem 5.1. In this connection, I think of the blocks of a partition of Σ (that is, of V) as *clusters* of vertices. (The name is meant to suggest cohesion within a block.) A *clustering* of a signed graph Σ is a partition π of Σ such that every edge within a cluster is positive and every edge between clusters is negative. Σ is called *clusterable* if it has a clustering, and *r-clusterable* if there is a clustering with at most r clusters.

Clearly, 2-clusterability is balance. The psychologist James A. Davis proved the analog of Harary's balance theorem:

Theorem 10.1 (Davis [9]). *Σ is clusterable if and only if no circle has exactly one negative edge.*

Clusterability differs from all other measures of imbalance, and most other properties of signed graphs, in that it is not invariant under switching. That is because signs of circles are not the essential property. This entails, for example, that we cannot treat $-K_P$ as an equivalent variant of K_P.

Write $E^-{:}\pi$ for the set of negative edges within clusters and $E^+(\pi)$ for the set of positive edges between clusters. These edges are *incompatible* with π. An *optimal partition* of Σ is a partition π of V such that Σ has the fewest possible edges that are incompatible with π. An optimal partition is a clustering if (and only if) Σ has any clustering at all. The cluster analog of frustration index is the *clusterability index*,

$$\operatorname{cli}(\Sigma) := \min_{\pi} |E^-{:}\pi| + |E^+(\pi)|, \text{ where } \operatorname{cli}_\pi(\Sigma) := |E^-{:}\pi| + |E^+(\pi)|$$

is the *inclusterability* of the partition π, i.e., the number of incompatible edges. (When $\pi = \{X, X^c\}$, this equals $|E^-(\Sigma^X)|$. To get a switching-like formula in general, write Σ^π for the signed graph obtained from Σ by negating all edges outside the clusters of the partition π; then $\operatorname{cli}_\pi(\Sigma)$, is the number of negative edges in Σ^π. The operation $\Sigma \mapsto \Sigma^\pi$ has no use I know of.) The clusterability index of Σ is the number of edges that are incompatible with an optimal partition, so it is a measure of how far Σ is from being clusterable. We could merely specify the number of clusters, defining

$$\operatorname{cli}_r(\Sigma) := \min_{\pi : |\pi| = r} \operatorname{cli}_\pi(\Sigma).$$

(The use of the clusterability and r-clusterability indices to find optimal clusterings is due to Doreian and Mrvar (1996a), who have a hill-climbing algorithm for probabilistic solution.)

The clusterability indices of $+P$ and $-P$ are easy to determine. The *cluster number* $\operatorname{cln}(\Sigma)$ is the minimum r such that Σ is r-clusterable, i.e.,

$$\operatorname{cln}(\Sigma) := \min\{r : \operatorname{cli}_r(\Sigma) = 0\}.$$

This is a finite number if and only if Σ is clusterable. Recall that $\pi(\Gamma)$ is the partition of V into the $c(\Gamma)$ vertex sets of the components of Γ. An *r-colored clustering* of Σ is a clustering whose clusters are colored by elements of an r-element set, so that clusters joined by a (negative) edge are colored differently.

Theorem 10.2. *The number of r-colored clusterings of a signed graph Σ is equal to the chromatic polynomial $\chi_{|\Sigma|/E^+}(r)$. Σ is clusterable if and only if the contraction $|\Sigma|/E^+$ has no negative loops.*

The clustering number $\operatorname{cln}(\Sigma) = \chi(|\Sigma|/E^+)$*, the chromatic number. Σ is r-clusterable if and only if $\chi(|\Sigma|/E^+) \leq r \leq c(\Sigma^+)$.*

An all-positive signed graph is r-clusterable if and only if $1 \leq r \leq c(\Gamma)$. For an all-negative signed graph, $\operatorname{cln}(-\Gamma) = \chi(\Gamma)$*, the chromatic number, and $-\Gamma$ is r-clusterable for* $\chi(\Gamma) \leq r \leq |V|$.

Proof. To get an r-colored clustering of Σ we simply have to color each component of Σ^+ with one of r colors, in such a way that components that are adjacent (negatively, of necessity) have different colors. This is the number of proper r-colorings of the graph which results from contracting the positive edges in Σ.

There exists a clustering $\iff$ there is a colored clustering for some number of colors $\iff$ $\chi_{|\Sigma|/E^+}(r)$ is not identically 0 $\iff$ $|\Sigma|/E^+$ has no negative loops. Furthermore, the smallest number of clusters in a clustering is the smallest number of colors in a proper coloring of the contracted graph. The largest number is $c(\Sigma^+)$ because a clustering of Σ partitions the components of Σ^+.

The all-positive and all-negative cases are easy deductions because $E^+(+\Gamma) = E$ and $E^+(-\Gamma) = \varnothing$. $\square$

Corollary 10.3. *$+P$, $-P$, $-P(e)$, and $-P_D$ are clusterable. $+P$ is r-clusterable only for $r = 1$; $-P$ is r-clusterable for $r \geq \chi(P) = 3$; $-P(e)$ is r-clusterable for $3 \leq r \leq 9$; and $-P_D$ is r-clusterable for $4 \leq r \leq 7$.*

Proof. Since $-P(e)$ and $-P_D$ are all negative except for a matching M_m of $m = 1$ or 3 edges, respectively, they are r-clusterable for $\chi(P/M) \leq r \leq 10 - m$. The chromatic numbers are $\chi(P/e) = 3$ and $\chi(P/M_3) = 4$, because M_3 consists of alternating edges of a hexagon. $\square$

Measurement of inclusterability.

With an inclusterable signed graph we can only look for a best approximation to a clustering. That has to be calculated in each case and there may be several optimal partitions.

There may be several optimal partitions, but they are constrained.

Theorem 10.4 (Doreian, Batagelj, Ferligoj, and Martin Everett [11, Theorem 10.6])**.** *The minimum clusterability index of Σ is attained on partitions of V whose sizes form a consecutive set of integers.*

The optimal partition sizes in Petersen examples may be one integer or several consecutive integers: see Corollary 10.3, Proposition 10.5, and Theorems 10.6 and 10.7.

Clustering of $P(e)$, or rather its failure, is easy to describe exactly. A less elementary example is P_D. Recall that the negative edges of P_D are alternating edges of a hexagon, which I will call the 'special hexagon' ($uvwxyzu$ in Figure 10.1(o)). The special hexagon has three chordal paths of length 2 that connect opposite vertices of the hexagon.

Proposition 10.5. *The clusterability index* $\mathrm{cli}(P(e)) = 1$*; the only optimal partition is the trivial one.*

Proof. Any partition of $P(e)$ other than the trivial one has at least 2 incompatible edges because P is edge 3-connected. ☐

Theorem 10.6. *The index* $\mathrm{cli}(P_D) = 3$*, from three types of partition: the trivial partition,* $\pi = \{X, X^c\}$ *where* X *is the vertex set of a negative edge, and* $\pi = \{Y, Y^c\}$ *where* Y *is the vertex set of a path* $uvwx$ *such that* uv *is a positive edge in the special hexagon and* vwx *is a chordal path.*

Proof. To show no partition π can have fewer than three incompatible edges we consider cases. Let $e = yz$, $f = uv$, $g = wx$ be the negative edges, where v and w are adjacent; see Figure 10.1(o).

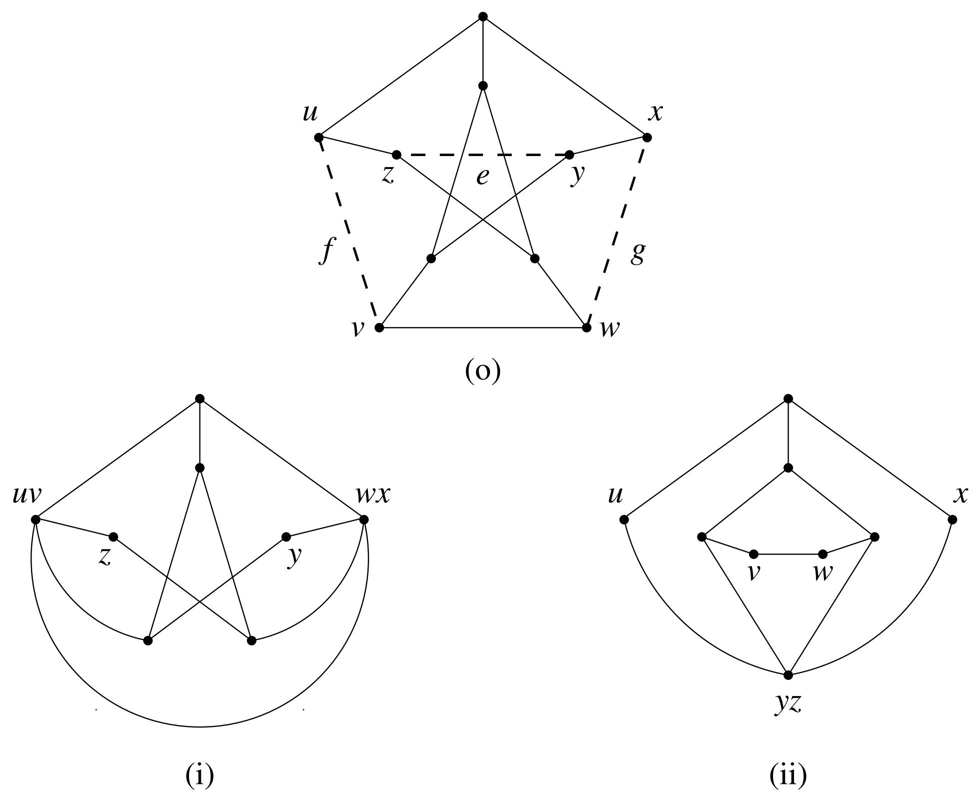

FIGURE 10.1. Cases 0–2 in the cluster analysis of P_D. (o) P_D with labels on the negative edges and their endpoints. (i) $P_D/fg \setminus e$. (ii) $P_D/e \setminus fg$.

The general method is to pick some negative edges to be within clusters, the others to be outside the clusters (interclustral edges). We contract the former and look at the edge connectivity between the endpoints of a negative interclustral edge to see whether only 3 edges can be incompatible with π. If a negative edge e connected two clusters, then any all-positive path between its endpoints must contain a positive interclustral edge, i.e., an incompatible positive edge. This lets us isolate the edges that must be incompatible, given the choice of interclustral negative edges.

Since P has automorphisms that permute e, f, g arbitrarily, every choice of a fixed number of edges to be within clusters behaves the same way. Thus, there are four cases, depending on how many negative edges are interclustral.

Case 0. All negative edges are within clusters. If π has more than one block, there are positive incompatible edges as well as the three negative ones. Thus, only the trivial partition is optimal.

Case 1. Two negative edges are within clusters, say f and g. Contracting f and g as shown in Figure 10.1(i), π is effectively partitioning $V(P/fg)$ so as to have e between clusters, i.e., its endpoints y and z should be in different clusters. These vertices are edge 2-connected within $(P/fg) \setminus e$, so there will be at least two incompatible positive edges; such a partition cannot be optimal.

Case 2. One negative edge is within a cluster, say e. Contracting e as in Figure 10.1(ii), we are effectively partitioning $V(P/e)$ so as to have f, g outside the clusters; that is, u, v are in different clusters and w, x are in different clusters. We study the structure of an optimal partition π.

In $P'' := P/e \setminus \{f, g\}$, u, v are joined by two edge-disjoint paths. Thus, to have them in different clusters there must be at least one edge from each path that is not in incompatible positive edges, we must pick two edges that separate both pairs and choose π so those edges are the only ones of P' outside clusters. The same two edges must separate v, w, since we are allowed only two positive edges between clusters. Consider the three paths $uv_ev'v$, $xv_ew'w$, and $utt'v'v$. The only way to disconnect all of them with two edges is for vv' to be one of the edges. Similarly, ww' must be one of the edges. Deleting these two separates V'' into $X := \{v, w\}$ and $V'' \setminus X$. Since each of them is connected in P'', π can only have two clusters, which are X and $V \setminus X$.

Case 3. All negative edges are between clusters. Then π partitions V so the endpoints of each negative edge are in different clusters; see Figure 10.2(a). Call vertices *equivalent* (relative to π) if they lie in the same cluster. Only three positive edges can be incompatible with π, i.e., interclustral. However, every uv, wx, or yz path in P_D^+ must contain an interclustral positive edge, since the endpoints are inequivalent. This is the basis of the solution.

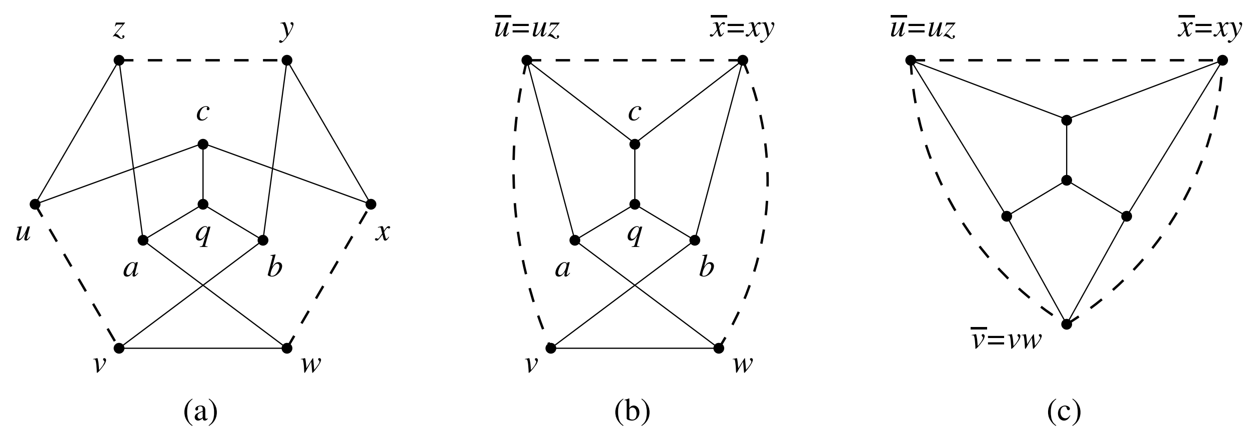

FIGURE 10.2. Case 3 in the cluster analysis of P_D. (a) The initial P_D. (b) Contracted by xy, uz; (c) and by vw.

At most one of the positive edges in the special hexagon, vw, xy, and uz, can be interclustral. If two of them were, say uz and vw, then in P_D^+ one edge would have to cut the paths $uaqbv$, $wcqax$, and $ybqcz$. However, no one edge is common to all these paths.

Thus, at least two positive edges in the special hexagon have equivalent endpoints; assume they are xy and uz. Contract them, as in Figure 10.2(b); let Q be the contraction. We need three edges whose deletion disconnects $\bar{x}$ from $\bar{u}$ and w and also $\bar{u}$ from v in Q^+.

If none of those edges is vw, we can contract vw to $\bar{v}$, as in Figure 10.2(c), and then it is impossible with three edges to separate every pair amongst $\bar{u}$, $\bar{v}$, and $\bar{x}$ in Q^+.

The only sets of three edges in Q^+ that include vw and separate the endpoints of every negative edge in Q are $\{vw, \bar{x}c, bq\}$ and $\{xy, \bar{u}c, aq\}$. In P_D these correspond to the sets $\{vw, xc, bq\}$ and $\{vw, uc, aq\}$, which are equivalent under symmetry of P_D. The optimal partition corresponding to $\{vw, xc, bq\}$ is a bipartition $\{Y, Y^c\}$ with $Y = \{v, x, y, b\}$, in other words, the vertices of an all-positive path of length 3 that starts with a positive edge of the special hexagon and ends at a vertex of that hexagon. This, up to symmetry, is the only optimal partition of P_D in which every negative edge is interclustral.

The conclusion is that no partition of P_D can have fewer than three incompatible edges, and we have found every optimal partition. □

Signed complete graphs are highly interesting.

Theorem 10.7. *Neither K_P nor $-K_P$ is clusterable.* $\mathrm{cli}(K_P) = 15$*, from just six partitions: the trivial partition, and the five partitions $\{X, X^c\}$ such that $X = N(u, w)$ for an edge $uw \in E(P)$.* $\mathrm{cli}(-K_P) = 10$*; an optimal partition of $-K_P$ is obtained by taking the clusters to be the endpoints of the edges of a perfect matching in P, and in no other way.*

Proof. Both graphs have obvious examples of circles with one negative edge; that proves they cannot be clustered.

The trivial partition, with one cluster, shows that K_P has clusterability index at most 15. This cannot be improved with two clusters, because of the frustration index. Suppose $\pi = \{X, X^c\}$. Then $(E^-{:}\pi) \cup E^+(\pi)$ is the set of edges incompatible with π. Switching X, this same set becomes the set of negative edges, which is no smaller than $l(K_P) = 15$ (Proposition 6.6). We know from Lemma 6.5 that we get as few as 15 negative edges only by switching $X = N(u, w)$ for a Petersen edge uw.

There is a similar analysis for partitions into more than two parts, but it cannot be stated in terms of switching. When $\Sigma = K_\Gamma$, the inclusterability of π is $\mathrm{cli}_\pi(\Sigma) = |E(\Gamma){:}\pi| + |E(\bar{\Gamma})(\pi)|$.

Lemma 10.8. *If a graph Γ of order n has the property that, for every $X \subseteq V$, the number of edges in the cut $E(\Gamma)(X, X^c)$ is $\leq \frac{1}{2}|X|(n - |X|)$, then* $\mathrm{cli}(K_\Gamma) = |E(\Gamma)|$.

Proof. In the preceding analysis of Σ, in the case $\Sigma = K_\Gamma$ the inclusterability is $\mathrm{cli}_\pi(K_\Gamma) = |E(\Gamma){:}\pi| + |E(\bar{\Gamma})(\pi)|$. The hypothesis implies that $|E(\Gamma)(X, X^c)| \leq |E(\bar{\Gamma})(X, X^c)|$. Hence, $\mathrm{cli}_\pi(K_\Gamma) \leq |E^-(K_\Gamma){:}\pi| + |E^-(K_\Gamma)(\pi)| = |E(\Gamma)|$. □

To treat K_P we only have to show that P satisfies the hypothesis of Lemma 10.8. Since every vertex has degree 3, this is easy for any cut except for a cut $E(K_P)(X, X^c)$ with $|X| = 5$. The number of Petersen edges contained within X is at least two, because either X has independence number ≤ 3, in which case it certainly contains at least two edges, or its independence number is 4, in which case it must be (up to symmetry) $\{12, 13, 14, 15\}$, and any additional vertex is joined to these four by at least two edges. Since X contains at least two Petersen edges, the cut $E(K_P)(X, X^c)$ contains at most $3|X| - 2(2) < \frac{1}{2}|E(K_P)(X, X^c)|$ Petersen edges. That solves K_P.

A partition that proves $\mathrm{cli}(-K_P) \leq 10$ is the partition given by a perfect matching, say

$$\pi_5 = \big\{\{12, 35\}, \{34, 25\}, \{15, 24\}, \{23, 14\}, \{45, 13\}\big\}.$$

Here $E^-{:}\pi_5 = \varnothing$ and $E^+(\pi_5)$ consists of the 10 (positive) edges of P that are not within the clusters. To analyze further we cannot use Lemma 10.8, since $\bar{P}$ plainly fails to satisfy the hypothesis. Instead, we analyze the contribution of each cluster to the total degree of the graph of incompatible edges.

That graph is $\Gamma_\pi := \big(V, (E^-{:}\pi) \cup E^+(\pi)\big)$. The number of incompatible edges is half the total degree of Γ_π. If every cluster contributes at least twice its own size to the total degree, then the number of incompatible edges is not less than $|V| = 10$. Let's consider a cluster X of k vertices. Within X there are m positive edges; that makes a contribution of $2\big[\binom{k}{2} - m\big]$ to the total degree of Γ_π. The number of incompatible (positive) edges departing X is $3k - 2m$, since P is cubic. The total contribution of X is therefore $k^2 + 2k - 4m$. Subtracting $2k$ we have $k^2 - 4m$, which we hope will always be nonnegative. If it is, then the only way to get as few as 10 incompatible edges is to have a partition with $k^2 - 4m = 0$ for every cluster.

Now the girth of P, namely 5, enters the picture. If $k \leq 4$, $m \leq k-1$, so $k^2 - 4m \geq (k-2)^2$, which is positive with one exception: when $k = 2$ and $m = 1$, i.e., when X is the vertex set of one Petersen edge. If $5 \leq k \leq 7$, $m \leq k$; then $k^2 - 4m \geq k(k-2) > 0$. If $k = 8$, the number of negative edges within the cluster is at least $\binom{8}{2} - m$, where $m \leq 10$ because at least 5 edges must be outside X; hence there are at least 18 incompatible edges within the cluster. When $k > 8$ there are only more negative incompatible edges. Thus, a cluster with $k \geq 8$ is itself more than enough to make $\mathrm{cli}_\pi > 10$.

Thus, any partition with a cluster of size other than 2, or with a cluster of size 2 that does not support a Petersen edge, has clusterability > 10. The optimal partitions are anything of the same form as π_5, and nothing else. All these partitions are equivalent under symmetries of P. $\square$

Part III. Algebra

11. The Positive Binary Cycle Subspace

From a certain algebraic viewpoint upon signatures of a graph the signed Petersen graphs are a peculiar family.

The characteristic vectors of edge sets in Γ form a vector space over $\mathbb{F}_2$, where we identify a set with its characteristic vector. The *binary cycle space*, $Z(\Gamma;\mathbb{F}_2)$, is the span of all circles, and it is also the set of all even-degree edge sets. The positive circles of a signature σ span a subspace $Z(\Sigma;\mathbb{F}_2)$, which we call the *positive binary cycle space*.

Theorem 11.1 (See, e.g., [29])**.** *Given a signed graph Σ with underlying graph Γ, a circle is positive in Σ if and only if it is in the positive binary cycle space $Z(\Sigma;\mathbb{F}_2)$. The codimension of $Z(\Sigma;\mathbb{F}_2)$ in $Z(\Gamma;\mathbb{F}_2)$ is 0 if Σ is balanced and 1 otherwise.*

Conversely, given a graph Γ, for any subspace Z of codimension 1 there is a signature σ such that $Z(\Sigma;\mathbb{F}_2) = Z$.

Outline of Proof. Reinterpret the signs $+1, -1$ as $0, 1 \in \mathbb{F}_2$. For the first half, apply the fact that, if the union of two circles is a theta graph, the sign of the third circle in the theta graph is the product of those of the first two.

For the second half, put positive signs on a maximal forest and sign any other edge e positive if its fundamental circle is in Z, negative if not. □

We draw two conclusions.

First, for a signed graph Σ on underlying graph Γ, $Z(\Gamma;\mathbb{F}_2)/Z(\Sigma;\mathbb{F}_2)$ is at most 1-dimensional. Therefore it can be regarded as $\subseteq \mathbb{F}_2$. The quotient map

$$Z(\Gamma;\mathbb{F}_2) \to Z(\Gamma;\mathbb{F}_2)/Z(\Sigma;\mathbb{F}_2) \hookrightarrow \mathbb{F}_2$$

is the sign function on circles. Thus, a switching class of signatures of Γ is equivalent to both a homomorphism $Z(\Gamma;\mathbb{F}_2) \to \mathbb{F}_2$ and a subspace Z of $Z(\Gamma;\mathbb{F}_2)$ whose codimension is at most 1.

Second, all circle signs are determined if we know those of circles that span $Z(\Gamma;\mathbb{F}_2)$. In particular, if the girth circles (the circles of minimum length in Γ) span the binary cycle space, their signs determine those of all circles. For instance, in K_n, $n \geq 3$, the signs of triangles determine those of all circles—which, truth to tell, is obvious from inspection. In particular, K_P is determined (up to switching) by its negative triangles.

In the Petersen graph, however, the pentagons do not span the binary cycle space. The dimension of the latter is 9, but:

Lemma 11.2. *The pentagons of P span a 5-dimensional subspace of $Z(P;\mathbb{F}_2)$.*

Proof. In the standard picture of P (Figure 1.1) there are four kinds of pentagons: the inner and outer pentagons C^I and C^O, five 'doubly outer pentagons' C^o_j with two edges in C^O, and five 'singly outer pentagons' with one edge in C^I.

We show that the doubly outer pentagons generate all others. Their sum is C^O. The sum of the singly outer pentagons is C^I. Figure 11.1 shows that C^O and the doubly outer pentagons generate the singly outer ones. We conclude that $\dim\langle\text{pentagons}\rangle \leq 5$. To prove the five doubly outer pentagons are linearly independent, notice that each contains one edge of C^I, which is not in any other doubly outer pentagon; thus there can be no dependencies amongst them. □

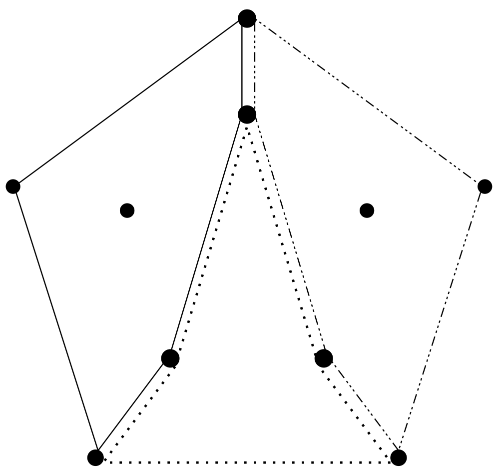

Figure 11.1. Pentagon dependencies in the Petersen graph. C^O is the sum of the three pentagons shown.

The lemma demonstrates that, after specifying the signs of all pentagons, there is still a 3-dimensional choice to be made to get an 8-dimensional positive cycle subspace Z. Therefore, there ought to be several signatures of P, not switching equivalent, for each choice of pentagon signs. Strangely, this is not what happens. Table 8.1 shows immediately that most choices of pentagon signs extend to circle signs in only one way, up to switching isomorphism. Only $P_{2,2}$ and P_D might conceivably have isomorphic pentagon signs and switching-nonisomorphic signatures. But, in fact, they do not.

Lemma 11.3. *There is no automorphism of P under which the set of negative pentagons of $P_{2,2}$ is carried to that of P_D.*

Proof. The pentagons of P come in vertex-disjoint pairs, such as C^O and C^I. One can check visually that the six negative pentagons in $P_{2,2}$ include two vertex-disjoint pairs but those of P_D do not include any vertex-disjoint pair. The proposition follows. □

This gives the main result, which is surprising in view of the preceding theoretical analysis.

Proposition 11.4. *A signature of P is determined up to switching isomorphism by its pentagon signs.*

Problem 11.1. Explain why the Petersen graph does not have more switching-nonisomorphic signatures for each pentagon signature. Is there a general reason that applies to other graphs? Is the key the difference between switching equivalence and switching isomorphism? Does Proposition 11.4 hold even up to switching equivalence?

12. Automorphisms [Incomplete]

An automorphism of a signed graph Σ is simply an automorphism of the underlying graph that preserves edge signs. This is not usually so interesting, because the focus on edge signs misses the main point of signed graphs, which is the circle signs, equivalently the switching class (cf. Proposition 4.1). Thus, I give most of my attention to *switching automorphisms*, which (naturally) are switching isomorphisms of Σ with itself. I write automorphisms on the right, as in Σ^α. The group identity is ε; for instance, in $\operatorname{Aut} P$ ε is the trivial permutation.

Automorphisms without switching.

Automorphisms of a signed graph without switching can be significant, e.g., when dealing with clusterability (Section 10), and they are generally easier to find as well. So we begin with them. $\mathrm{Aut}(\Sigma)$ denotes the automorphism group of Σ. The first, and very obvious, result takes care of most of our Petersen examples.

Proposition 12.1. *For any signed graph,* $\mathrm{Aut}(\Sigma) = \mathrm{Aut}(-\Sigma) = \mathrm{Aut}(\Sigma^+) \cap \mathrm{Aut}(\Sigma^-)$.
For any graph, $\mathrm{Aut}(+\Gamma) = \mathrm{Aut}(-\Gamma) = \mathrm{Aut}(\Gamma)$.
For a simple graph, $\mathrm{Aut}(K_\Gamma) = \mathrm{Aut}(-K_\Gamma) = \mathrm{Aut}(\Gamma)$. □

An inhomogeneously signed graph is rather different from a homogeneous one; Σ may have many fewer automorphisms than $|\Sigma|$. That is certainly the case with $P(e)$ and P_D. The task is to find the exact automorphism group. The symmetric group of degree k is $\mathfrak{S}_k$. The full permutation group of a set X is $\mathfrak{S}_X$.

Proposition 12.2. $\mathrm{Aut}\, P(e) \cong \mathfrak{S}_2 \wr \mathfrak{S}_2$ *(the wreath product) and* $\mathrm{Aut}\, P_D \cong \mathfrak{S}_3$.

Lemma 12.3. *Let H be a hexagon in P. Any automorphism of H extends uniquely to an automorphism of P.*

Proof. Let α_H be the automorphism of H. Fix a 4-arc A (a directed path of length 3) in H. Clearly, α_H is determined by its restriction to A. There is a unique automorphism α of P that moves A to A^{α_H}. This automorphism preserves H (that is, $H^\alpha = H$) because any 4-arc lies in a unique hexagon. Therefore, α is the required unique extension. □

Proof of Proposition 12.2. First consider $P(e)$ with negative edge $e := \{12, 34\}$. An automorphism of $P(e)$ is simply an automorphism of P that fixes e. (It may interchange the endpoints.) Thus, it is a permutation of $[5]$ that fixes 5 and can independently exchange $1 \leftrightarrow 2$, $3 \leftrightarrow 4$, and $\{1, 2\} \leftrightarrow \{3, 4\}$.

More precisely, if the automorphism α fixes the vertices 12 and 34, it may exchange 35 and 45 and, independently, 15 and 25. Thus, it belongs to the group $\mathfrak{G} := \mathfrak{S}_{\{3,4\}} \times \mathfrak{S}_{\{1,2\}}$. An automorphism that exchanges 12 and 34, such as $\psi := (13)(24)$, belongs to the coset $\psi\mathfrak{G}$. The whole of $\mathrm{Aut}\, P(e)$ therefore is the semidirect product of $\langle\psi\rangle$ and $\mathfrak{G}$, i.e., $\mathfrak{S}_2 \wr \mathfrak{S}_2$.

In P_D, the three negative edges e_1, e_2, e_3 are alternating edges of a hexagon H. Any permutation of E^- determines a unique automorphism of H, which by the lemma extends uniquely to an automorphism of P, which clearly belongs to $\mathrm{Aut}\, P_D$. Therefore, $\mathrm{Aut}\, P_D$ contains Σ_{E^-}. On the other hand, an automorphism α of P_D must preserve E^-; thus, it is determined by its action on E^-. We conclude that $\mathrm{Aut}\, P_D = \Sigma_{E^-} \cong \mathfrak{S}_3$. □

Switching automorphisms.

When we come to switching automorphisms the picture is quite different. First we have to define one precisely. I will give a definition that is suitable to simple graphs like P.

A *switching permutation* of V is a pair (τ, α) of a switching function $\tau \in \{+1, -1\}^V$ and $\alpha \in \mathfrak{S}_V$. Its action on Σ is defined as the permutation α applied to the switched graph Σ, i.e., as $\Sigma^{\tau\alpha}$. (Thus, I will write $\tau\alpha := (\tau, \alpha)$ from now on.) This gives a switching isomorphic graph in the sense of Section 4: isomorphic as a signed graph to a switching of Σ. When $\Sigma^{\tau\alpha}$ is Σ itself then we say the action of $\tau\alpha$ is a *switching automorphism* of Σ. Let $\mathrm{SwPerm} := \{\tau\alpha : \Sigma^{\tau\alpha} = \Sigma\}$.

These are not yet switching automorphisms. The necessary and sufficient condition for $\Sigma^\tau = \Sigma$ is that τ be constant on each component of Σ. Thus, the subgroup $\mathfrak{K} := \{\tau\varepsilon_V :$

$\Sigma^\tau = \Sigma\} \leq \mathrm{SwPerm}(\Sigma)$ has a trivial action on Σ. The group of switching automorphisms of Σ is therefore the quotient $\mathrm{SwAut}\,\Sigma := \{(\mathrm{SwPerm}\,\Sigma)/\mathfrak{K}$. In concrete terms, a switching automorphism is a coset of $\mathfrak{K}$ in $\mathrm{SwPerm}\,\Sigma$. Perhaps even more concretely, it is a coset of $\mathfrak{K}$ in $\{+1,-1\}^V \times \mathfrak{S}_V$, acting on the class of signed simple graphs on vertex set V, which fixes Σ.

Despite this technical discussion I shall often call $\tau\alpha$ a switching automorphism when it is really the coset $\tau\alpha\mathfrak{K}$ that is the switching automorphism.

The significance of switching automorphisms concerns circle signs (cf. Proposition 4.2).

Proposition 12.4. *If $\tau\alpha$ is a switching automorphism of Σ, then α is an automorphism of $|\Sigma|$ that preserves the signs of circles. Conversely, if α is such an automorphism of $|\Sigma|$, there is a switching function τ such that $\tau\alpha$ is a switching automorphism of Σ, and τ is unique up to negation on components of Σ.*

The automorphisms of Σ form a subgroup of the switching automorphisms, but it need not be normal (cf. K_P, for instance). No nontrivial switching is a switching automorphism, as long as the underlying graph is simple. (With multiple edges of opposite sign, a switching can be a nontrivial automorphism; but that does not apply to any Petersen signed graph.) Half the work is saved by noticing the simple behavior of negation:

Proposition 12.5. $\mathrm{SwAut}(-\Sigma) = \mathrm{SwAut}(\Sigma)$. □

For a homogeneous signed graph like $+P$ or $-P$ there is nothing new here. There can be no nontrivial switching, with or without an automorphism of the signed graph, because the result would not be homogeneous with the same sign.

Proposition 12.6. $\mathrm{SwAut}(+\Gamma) = \mathrm{SwAut}(-\Gamma) = \mathrm{Aut}\,\Gamma$. □

However, as soon as Σ is inhomogeneous there may be switching automorphisms that are not automorphisms—or, there may not.

Proposition 12.7. $\mathrm{SwAut}\,P(e) = \mathrm{Aut}\,P$.

Proof. In $P(e)$, nontrivial switching makes more than one negative edge because no cut has fewer than three edges. Hence, if τ is nontrivial, $P(e)^{\tau\alpha}$ cannot be $P(e)$. □

Switching automorphisms of P_D.

Theorem 12.8. $\mathrm{SwAut}\,P_D =$???.

Proof. □

Switching automorphisms of K_P.

Now we move on to K_P. Let's set up some notation. The permutation of $V := V(P) = \mathcal{P}_2([5])$ induced by a permutation ψ of [5] is α_ψ; it is an automorphism of P. In particular, α_{ij} is induced by the transposition $(i\ j)$. Another permutation of V is defined directly on V. Let i, j, k, l, m be the five elements of [5] in any order. Then

$$\gamma_i := (jk\ lm)(jl\ km)(jm\ kl),$$

the product of transpositions on V.

Each maximum independent set has the form $X_i := \{ij : j \in [5] \setminus i\}$. Define $\tau_i := \tau_{X_i}$ to be the function that switches X_i, i.e., $\tau_i^{-1}(-1) = X_i$. Now, let

$$\beta_i := \tau_i \gamma_i \qquad \text{for } i = 1, \ldots, 5,$$
$$\beta_0 := \varepsilon \qquad \text{for } i = 0.$$

These functions are switching automorphisms of K_P, as we have already seen less explicitly in Lemma 6.5.

Lemma 12.9. *Let $X \subseteq V(K_P)$. $(K_P)^X$ is isomorphic to K_P if and only if $X = \varnothing$ or X_i for some $i = 1, 2, \ldots, 5$. When $X = X_i$ the isomorphism is any element of the coset $\gamma_i \operatorname{Aut} P$ in the symmetric group on V.*

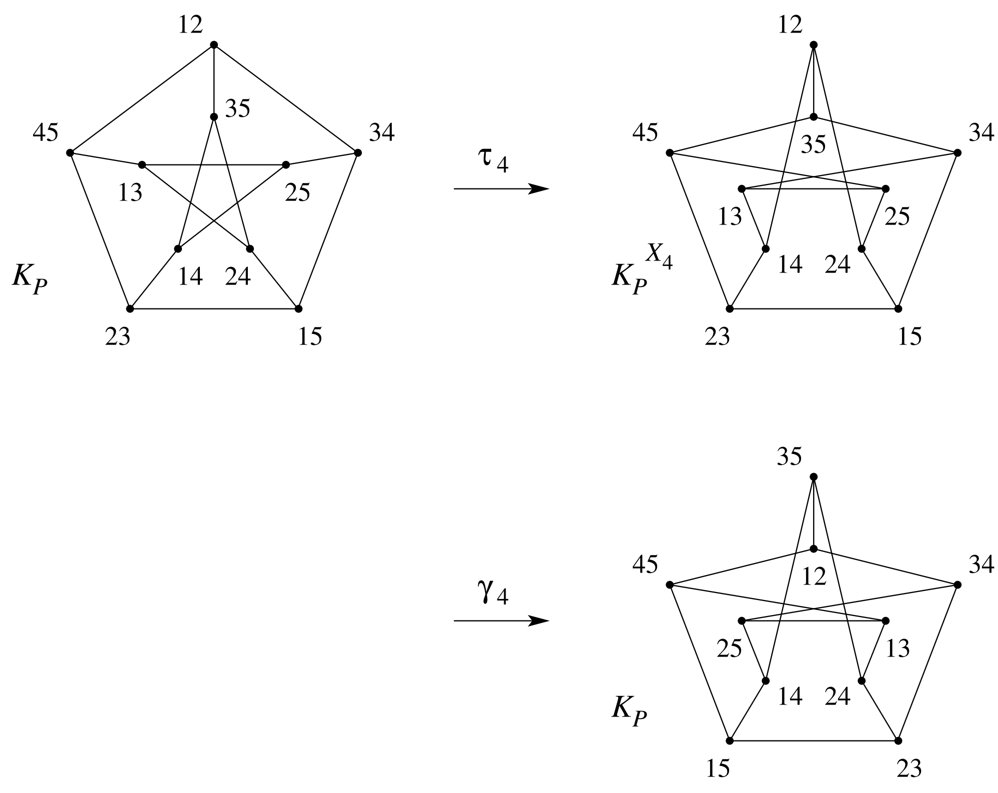

FIGURE 12.1. The action of $\beta_i = \tau_i \gamma_i$ in two stages, for $i = 4$.

Proof. Comparing with the proof of Lemma 6.5, the only thing lacking is that γ_i itself is an isomorphism $(K_P)^{\tau_i} \to K_P$, which is easy enough to check. $\square$

Now define $\theta : \mathfrak{S}_6 \to \operatorname{SwAut} K_P$ by extending

$$\theta\big((i\ j)\big) = \alpha_{ij} \qquad \text{for a transposition } (i\ j) \in \mathfrak{S}_5,$$
$$\theta\big((i\ 6)\big) = \beta_i \qquad \text{for } i \in [5].$$

Theorem 12.10. *The group* $\operatorname{SwAut} K_P$ *is isomorphic to* $\mathfrak{S}_6$ *under the mapping* θ*. The subgroup* $\operatorname{Aut} K_P$ *corresponds to* $\mathfrak{S}_5 \leq \mathfrak{S}_6$*.*

The switching automorphisms of K_P *have the form* $\beta_i\alpha$ *where* $i = 0, 1, \ldots, 5$ *and* $\alpha \in \operatorname{Aut} P$*.*

We need to prove θ is well defined, one-to-one, and onto. All this is done by combinatorial group formulas. Thus, our first task is to set up the elementary relations amongst the α's, β's, γ's, and τ's. A reader who isn't entertained by the elementary algebraic manipulations involved in the set-up can prove Equations (12.2) from the definitions, if desired, or go directly to the proof of Theorem 12.10.

The important operations for the group description are the α_{ij} and β_i. The tools for working with them are the γ_i and τ_i. Thus, we begin with relations amongst the latter two and the α_{ij}.

Remember that $\tau_{\{ij\}}$ is the function that switches the vertex ij. Assume that i, j, k, l, m are distinct elements of [5] unless stated otherwise.

$$\gamma_i^2 = \varepsilon\,; \tag{12.1a}$$

$$\begin{aligned} \tau_{\{ij\}}\gamma_m &= \gamma_m\tau_{\{kl\}}, \\ \tau_{\{ij\}}\gamma_i &= \gamma_i\tau_{\{ij\}}\,; \end{aligned} \tag{12.1b}$$

$$\alpha_{ij} = \gamma_i\gamma_j\gamma_i = \gamma_j\gamma_i\gamma_j\,; \tag{12.1c}$$

$$\gamma_i\gamma_j = \alpha_{ij}\gamma_i = \gamma_j\alpha_{ij}\,; \tag{12.1d}$$

$$\begin{aligned} \gamma_i\alpha_{ij} &= \alpha_{ij}\gamma_j, \\ \gamma_i\alpha_{kl} &= \alpha_{kl}\gamma_i\,; \end{aligned} \tag{12.1e}$$

$$\begin{aligned} \gamma_i\tau_i &= \tau_i\gamma_i, \\ \gamma_i\tau_j &= \tau_j\tau_i\gamma_i. \end{aligned} \tag{12.1f}$$

These have easy direct proofs using cycle or 1- or 2-line forms of the permutations involved, but I want to show how one can set up a calculus that avoids having to work with such explicit details.

Equation (12.1a) is obvious from the definition. Equation (12.1b) is a most basic identity, proved by

$$\tau_{\{ij\}}\gamma_m = \tau_{\{ij\}}(ij\ kl)(ik\ jl)(il\ jk) = (ij\ kl)(ik\ jl)(il\ jk)\tau_{\{kl\}} = \gamma_m\tau_{\{kl\}}$$

and

$$\tau_{\{ij\}}\gamma_i = \tau_{\{ij\}}(jk\ lm)(jl\ km)(jm\ kl) = (jk\ lm)(jl\ km)(jm\ kl)\tau_{\{ij\}} = \gamma_i\tau_{\{ij\}}.$$

Equation (12.1c) is also most basic; its proof is another easy direct calculation:

$$\begin{aligned} \gamma_i\gamma_j\gamma_i &= (kl\ mj)(km\ lj)(kj\ lm)(kl\ mi)(km\ li)(ki\ lm)(kl\ mj)(km\ lj)(kj\ lm) \\ &= (ki\ kj)(li\ lj)(mi\ mj)(ij) = \alpha_{ij}. \end{aligned}$$

From now on I apply these tools as needed to prove other formulas.

For instance, (12.1d) is simply a rearrangement of (12.1c) using (12.1a). The first case of (12.1e) is a rearrangement of (12.1d), but the second case involves (12.1c) and (12.1d):

$$\gamma_i\alpha_{kl} = \gamma_i\gamma_k\gamma_l\gamma_k = \alpha_{ik}\gamma_i\gamma_l\gamma_k = \cdots = \alpha_{ik}\alpha_{il}\alpha_{ik}\gamma_i = \alpha_{kl}\gamma_i$$

because of properties of transpositions, which satisfy $(i\ k)(i\ l)(i\ k) = (k\ l)$.

The first case of (12.1f) holds true because γ_i leaves X_i fixed. In the second case I employ the fact that $\tau_X\tau_Y = \tau_{X\oplus Y}$ (where $\oplus$ denotes set summation):

$$\begin{aligned}
\gamma_i\tau_j &= \gamma_i\tau_{\{ij\}}\tau_{\{jk\}}\tau_{\{jl\}}\tau_{\{jm\}} \\
&= \tau_{\{ij\}}\tau_{\{lm\}}\tau_{\{km\}}\tau_{\{kl\}}\gamma_i \\
&= \tau_j\tau_{\{jk,jl,jm,kl,km,lm\}}\gamma_i \\
&= \tau_j\tau_i\gamma_i
\end{aligned}$$

because $X_i = \{jk, jl, jm, kl, km, lm\}^c$.

From Equations (12.1) we deduce that

$$\begin{aligned}
\tau_i\alpha_{ij} &= \alpha_{ij}\tau_j, \\
\tau_k\alpha_{ij} &= \alpha_{ij}\tau_k\,;
\end{aligned} \tag{12.2a}$$

$$\beta_i^2 = \varepsilon\,; \tag{12.2b}$$

$$\beta_i\beta_j\beta_i = \beta_j\beta_i\beta_i = \alpha_{ij}\,; \tag{12.2c}$$

$$\begin{aligned}
\beta_i\alpha_{ij} &= \alpha_{ij}\beta_j, \\
\beta_i\alpha_{kl} &= \alpha_{kl}\beta_i.
\end{aligned} \tag{12.2d}$$

For instance,

$$\beta_i^2 = \tau_i\gamma_i\tau_i\gamma_i = \tau_i^2\gamma_i^2 = \varepsilon$$

from (12.1f) in particular. For (12.2a), extract the τ's and apply (12.1c):

$$\alpha_{ij}\tau_i = \gamma_i\gamma_j(\gamma_i\tau_i) = \gamma_i(\gamma_j\tau_i)\gamma_i = (\gamma_i\tau_i\tau_j)\gamma_j\gamma_i = \tau_j(\gamma_i\gamma_j\gamma_i) = \tau_j\alpha_{ij}$$

and

$$\alpha_{ij}\tau_k = \gamma_i\gamma_j(\gamma_i\tau_k) = \gamma_i(\gamma_j\tau_k\tau_i)\gamma_i = (\gamma_i\tau_k\tau_i)\gamma_j\gamma_i = \tau_k(\gamma_i\gamma_j\gamma_i) = \tau_k\alpha_{ij}.$$

For (12.2c), also move the τ's around and apply (12.1c):

$$\beta_i\beta_j\beta_i = \tau_i(\gamma_i\tau_j\gamma_j\gamma_i)\tau_i = (\tau_j\gamma_i)\gamma_j\gamma_i\tau_i = \tau_j\alpha_{ij}\tau_i = \alpha_{ij}$$

by (12.2a). For (12.2d) employ the commutation relations of (12.1d) and (12.2a).

Proof of Theorem 12.10. For well definition of θ we show that $\beta_i \in \operatorname{SwAut} K_P$ and $(i\ 6) \in \mathfrak{S}_6$ have the same order and the same commutation relations with the corresponding transpositional elements of their groups, α_{ij} and $(i\ j)$. For these relations see, e.g., [19, ???].

The order of β_i is given by (12.2b). The commutation relations contained in (12.2c, d) are the same as the fundamental relations of the transpositions $(i\ 6)$. It follows that θ is well defined and a homomorphism, hence an epimorphism.

To show θ is injective it suffices to prove that all β_i are different. This is obvious: each one both switches and permutes V differently. $\square$

There is another, relatively simple way to prove that $\mathfrak{S}_6$ is the automorphism group of K_P. The construction of K_P as the quotient of the Johnson graph $J(6,3)$ modulo the complementation involution (Section 17) implies it. This proof is in Section **??**.

Part IV. Matrices

13. Adjacency Matrix

The *adjacency matrix* of a signed simple graph, $A(\Sigma)$, is the $V \times V$ matrix whose entry $a_{ij} = \sigma(e_{ij})$ if v_i and v_j are adjacent and 0 if they are not. More generally, if Σ has parallel edges then a_{ij} is the sum of the signs of all edges v_iv_j. For instance, $A(+\Gamma) = A(\Gamma)$ and $A(-\Gamma) = -A(\Gamma)$. The matrices of graphs K_Γ are almost as simple:

$$A(K_\Gamma) = J - I - 2A(\Gamma),$$

where J is the all-ones matrix. It will be no surprise, then, that the matrix and eigenvalue properties of signed graphs of any of these types are intimately related to those of Γ.

Obvious properties are that $A(\Sigma)$ is a symmetric, integral matrix with zero diagonal. Powers of A count walks of given length between vertices, but here there is a modification: the (i, j) entry of A^l is the number of walks of length l from v_i to v_j that are positive, less the number that are negative. This can easily equal 0 even when walks exist.

Switching has a nicely simple effect on the adjacency matrix. Suppose $\tau : V \to \{+1, -1\}$ is a switching function. Let $D(\tau)$ be the diagonal matrices with entries given by τ. Then

$$A(\Sigma^\tau) = D(\tau)^{-1}A(\Sigma)D(\tau). \tag{13.1}$$

(Since $D(\tau)^{-1} = D(\tau)$, the inversion is merely ideological, to show that A is being conjugated.)

Many properties of $A(\Sigma)$ resemble those of an unsigned graph, but with little differences. For instance, powers A^l count walks of length l between pairs of vertices, though the count is distorted by the fact that positive and negative walks cancel. The diagonal entries of A^2, if A has no loops, are the vertex degrees in $|\Sigma|$. The diagonal entries of A^3 get interesting. If $|\Sigma|$ has girth > 3, as does P, the diagonal is zero. If $\Sigma = K_\Gamma$, the diagonal entry of vertex v is the number of positive triangles on v less the number of negative ones; switching so all edges at v are positive, it is the number of negative edges in the switched graph. We say more about signed complete graphs in Section **??**.

14. Incidence Matrix

A graph has two common incidence matrices. The unsigned incidence matrix has two $+1$'s in each column, and the oriented incidence matrix has one $+1$ and one -1. Both are examples of the incidence matrix of a signed graph.

The *incidence matrix* of Σ is a $V \times E$ matrix $\mathrm{H}(\Sigma) = (\eta_{ve})_{v,e}$ (read H as "Eta") in which the column of an edge e_{ij} has ± 1 in the row of v_i and $\pm\sigma(e_{ij})$ in the row of v_j, and is 0 elsewhere. (The $\pm$ must be the same in both rows.) Obviously, this is not unique, since negating a column gives a different matrix for Σ, but the difference, which is a matter of choice of orientation, will not matter for the time being. This definition applies to edges that are not loops. The column of a positive loop is all 0 and that of negative loop e_{ii} is 0 except for a ± 2 in the row of v_i.

As examples, the incidence matrix of $+\Gamma$ is the oriented incidence matrix of Γ and $\mathrm{H}(-\Gamma)$ is (after choosing the column signs suitably) the unsigned incidence matrix of Γ.

The incidence matrix satisfies some of the usual properties. Let $\Delta(\Gamma)$ be the diagonal degree matrix of Γ. (A loop counts twice in the degree.)

Theorem 14.1. *For a signed graph,* $\mathrm{H}(\Sigma)\mathrm{H}(\Sigma)^{\mathrm{T}} = \Delta(|\Sigma|) - A(\Sigma)$*. If* $|\Sigma|$ *is* k*-regular, then* $A(\Sigma) = kI - \mathrm{H}(\Sigma)\mathrm{H}(\Sigma)^{\mathrm{T}}$.

The proof is standard: compare the two sides, element by element. The matrix $K(\Sigma) := \Delta(|\Sigma|) - A(\Sigma)$ is the *Kirchhoff matrix* of Σ. (It is often called the *Laplacian matrix* but other matrices are also called that.) For an unsigned graph, the theorem says that

$$\begin{aligned} A(\Gamma) &= A(+\Gamma) = \Delta(\Gamma) - \mathrm{H}(+\Gamma)\mathrm{H}(+\Gamma)^{\mathrm{T}}, \\ A(\Gamma) &= -A(-\Gamma) = \mathrm{H}(-\Gamma)\mathrm{H}(-\Gamma)^{\mathrm{T}} - \Delta(\Gamma). \end{aligned} \tag{14.1}$$

Thus we have two standard graphical formulas in one package.

The columns of the incidence matrix give a vector representation of Σ whose linear dependencies can be described in strictly graphical terms. This is matroid theory, which I omit. Here I observe that the column vectors belong to a well-known set: a root system C_n. Rather than giving a general definition I will state the canonical form of each type we need here. The standard orthonormal basis of $\mathbb{R}^n$ is the set $\{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$. These are three of the "classical" root systems:

A_{n-1} is the set of all vectors in $\mathbb{R}^n$ of the form $\mathbf{b}_j - \mathbf{b}_i$ where $i \neq j$.
D_n is $A_{n-1} \cup \{\pm(\mathbf{b}_j + \mathbf{b}_i) : i \neq j\}$.
C_n is $D_n \cup \{\pm 2\mathbf{b}_i\}$.

The column of a positive non-loop edge belongs to A_{n-1}, that of any non-loop to D_n, and that of any edge except a positive loop to C_n. The zero vector, corresponding to a positive loop, does not belong to any root system.

Recall that $b(\Sigma)$ is the number of balanced components of Σ. A main theorem of signed graph theory is:

Theorem 14.2 ([30, Theorem 8A.1])**.** *The rank of* $\mathrm{H}(\Sigma)$ *equals* $n - b(\Sigma)$.

Thus, all the incidence matrices of Petersen signed graphs have rank 10, except for $+P$, where the rank is 9 since the graph is balanced.

Part V. Covering Graphs and Surface Embedding

15. Signed Covering Graphs

The *signed covering graph* of Σ (also known as the *derived graph*; cf. Biggs (1974a), Exercise 19A) is the unsigned graph $\widetilde{\Sigma}$ whose vertex set is $\widetilde{V} := V \times \{+1, -1\}$ with edges $\widetilde{e}_{ij}$, joining $(v_i, +1)$ to $(v_j, +\sigma(e_{ij}))$, and $\widetilde{e}^*_{ij}$, joining $(v_i, -1)$ to $(v_j, -\sigma(e_{ij}))$. The mapping * defined by $(v_i, \varepsilon)^* := (v_i, -\varepsilon)$ and $(\widetilde{e}^*_{ij})^* := \widetilde{e}_{ij}$ is an automorphism of $\widetilde{\Sigma}$ of period 2, with no fixed points. The *covering projection* $p : \widetilde{\Sigma} \to \Sigma$ defined by $p(v_i, \varepsilon) = v_i$ and $p(\widetilde{e}_{ij}) = p(\widetilde{e}^*_{ij}) := e_{ij}$ is a graph homomorphism. If we forget the vertex signs in $\widetilde{\Sigma}$ we have a *double covering graph* of Σ that we write $\widetilde{\Gamma}$; then * is a fixed-point-free, involutory automorphism of $\widetilde{\Gamma}$.

Conversely, if $\widetilde{\Gamma}$ is any graph with a fixed-point-free, involutory automorphism *, then $\widetilde{\Gamma}/^*$ is a signed graph. Actually, that is a bit of an overstatement. In order to get a signed graph we have to choose one of each pair of corresponding vertices of $\widetilde{\Gamma}$, v and v^*, to be called positive and call the other one negative. Then we can define the sign of an edge in the quotient graph, i.e., $\sigma(\{\widetilde{e}, \widetilde{e}^*\})$, to be the product of the signs of the endpoints of $\widetilde{e}$ (or of $\widetilde{e}^*$; it's the same sign). This construction gives us a signed graph Σ of which $\widetilde{\Gamma}$ is the double covering graph (if we are not too pedantic about notation).

This construction in different terminology is well known. See for instance [2, Problem 19A]. It is a special case of constructions of interest to group theorists [2, Chapter 19] and topological graph theorists (cf. Gross (1974a)). Our interest is more specialized; but these two topics will be the subject of Sections **??** and 16.

Changing the choice of how to sign the vertex pairs in $\widetilde{\Gamma}$ corresponds to switching the changed vertices in Σ. Thus, $\widetilde{\Gamma}/^*$ is really not a signed graph but a switching equivalence class. The reason is that, in the signed covering, we identify vertices within a covering pair as positive and negative, but in the double covering we do not; if we supply this additional information along with $\widetilde{\Gamma}$ the projection is a unique signed graph.

The natural next question is: What are the signed covering graphs of various signatures of P? I will answer this as well as I can in the following sections. The discussion shows that a switching class of signed graphs is equivalent to an ordinary graph together with an involutory automorphism without fixed vertices, and a signed graph is equivalent to an ordinary graph, a fixed-vertex-free involutory automorphism, and a sign labelling of pairs of automorphically equivalent vertices. We have two uses for this observation: one in relation to surface embedding (Section 16) and one directly concerning automorphisms (Section 17). Each approach demonstrates a double covering of a Petersen signed graph.

There is one special case to dispose of at once. If Σ is balanced, $\widetilde{\Sigma}$ consists of two identical copies of $|\Sigma|$ and * is the natural bijection between the copies. The proof is that Σ switches to all positive, and then $\widetilde{\Sigma}$ has one copy of $|\Sigma|$ on the positive vertices and another on the negative ones. This is not very interesting. In particular, $\widetilde{+P}$ consists of two copies of P with an isomorphism $\alpha : (v, +) \mapsto (v, -)$ serving as the involution. A simple special case is important:

Lemma 15.1. *Let C be a circle in Σ. Then $p^{-1}(C)$ is one circle if C is negative and it is two circles, each isomorphic to C, if C is positive.* □

A more interesting special case is an antibalanced Σ. The signed covering is bipartite, as one can see by switching Σ to be all negative and observing that then every edge in $\widetilde{\Sigma}$ has one positive endpoint and one negative endpoint. Such graphs can be quite interesting. $-P$ is a case in point. Its double covering graph is the Desargues graph, which is the point-line incidence graph of the Desargues configuration from projective geometry. It is a cubic distance-regular graph of order 20, bipartite, with girth 6, and it is the Cartesian product $K_2 \square P$. (How we know this might require explanation. A.E. Brouwer [3] observes that the Desargues graph is the "bipartite double" of P, which means $K_2 \square P$. This in turn is simply $\widetilde{-P}$, if we take the vertex set of K_2 to be $\{+1, -1\}$ and compare definitions.)

The signed cover of a graph with one negative edge, like $\widetilde{P}(e)$, is easy to describe. Suppose the negative edge of $P(e)$ is u_1w_1. In $\widetilde{+P}$, replace u_1w_1 in the first copy and u_2w_2 in the second copy by edges u_1w_2 and u_2w_1. The involution of this graph is α. This gives $\widetilde{P}(e)$.

We shall see in Section 16 that the double covering of P_D is the dodecahedral graph D—whence the notation for this signature.

Section 17 will show what is interesting about $\widetilde{K}_P$.

16. Surface Embedding

There is more than one way to define surface embedding of a signed graph, but that which seems to me the most natural is what I call orientation embedding. The basic ideas here are due, independently, to Gerhart Ringel [18], Saul Stahl [25], and me ([31], on which I mainly rely here). (Gross and Tucker's book (1987a) has extensive treatment of orientation embedding in different terminology.)

We need notation for the closed surfaces; I call them T_g for the g-fold torus, the orientable surface with g handles ($g \geq 0$) and U_h for the nonorientable surface with h crosscaps. The sphere is T_0, the torus is T_1, the projective plane is U_1, the Klein bottle is U_2; these are the most fun. The most important fact to know is that T_g has an involutory self-homeomorphism under which it is the unique orientable double covering surface of U_h, where $g = 2h - 2$; for example, the unit sphere has the antipodal mapping given by its central symmetry, and the quotient surface is U_1. Thereby any graph $\widetilde{\Gamma}$ embedded in T_{2h-2} in conformity with the involution projects to a signed graph Σ in U_h, and by comparing the definitions one can see that $\widetilde{\Gamma}$ is the double covering of Σ.

A signed graph Σ is *orientation embedded* in a surface S if its underlying graph is topologically embedded in S so that a positive circle preserves orientation and a negative circle reverses it. Obviously, Σ embeds in some orientable surface if and only if it is balanced; consequently, orientation embedding of a balanced signed graph is the same as orientable-surface embedding of the underlying graph. There is always a unique minimal surface $S(\Sigma)$ in which Σ can be orientation embedded; this surface is orientable if and only if Σ is balanced. The *demigenus* $d(S)$ of a surface (also called *Euler genus* or *nonorientable genus*) is defined as $2 -$ Euler characteristic). The *demigenus of* Σ, $d(\Sigma)$, is the demigenus of its minimal surface. Thus, the minimal surface of Σ is uniquely determined if we know two things about Σ: its demigenus and whether it is balanced or not.

Each nonorientable topological embedding of $|\Sigma|$ determines circle signs, hence a switching class of signed graphs (Proposition 4.1). There is a simple and convenient way to derive edge signs from a topological embedding: draw the embedding in a polygonal diagram of the surface with boundary identification scheme $a_1a_1a_2a_2\cdots a_h$, with no vertices on the boundary. The sign of e is $(-1)^\kappa$ where κ is the number of times e crosses the boundary.

Let's apply all this to signed Petersen graphs.

Proposition 16.1. *The demigenera of some signed Petersen graphs are $d(+P) = 2$, $d(P_D) = 1$, $d(-P) = 4$, and $d(P(e)) = 3$.*

Proof. $+P$ is obviously nonplanar (who can help noticing how it contracts to K_5?) but it does embed in the torus (see [15, Figure 9.5]) so $d(+P) = 2$. (The fact that P embeds in the projective plane as an unsigned graph is immaterial to $+P$ because all circles in the latter must be orientation preserving, which forces the minimal embedding surface to be orientable.)

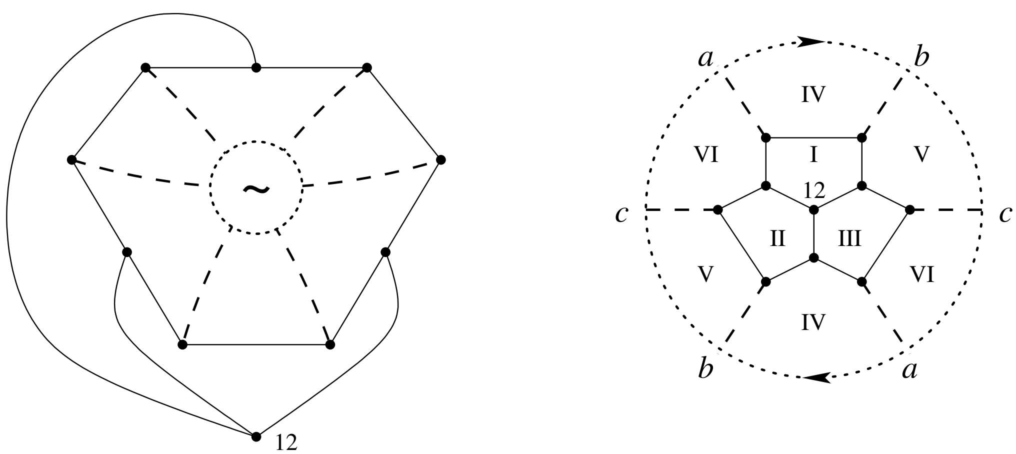

FIGURE 16.1. Two views of P, signed as P_D, embedded in the projective plane. Left: The central circle and twiddle represent a crosscap; the edges through the crosscap are the negative ones. Right: The outer circle represents the crosscap; the six pentagonal regions show duality with K_6.

The unsigned Petersen graph embeds in the projective plane U_1 (see Figure 16.1), uniquely because it is the surface dual graph of K_6, whose embedding in U_1 is unique (cf. [15, Section 9.9]). This embedding gives the signs of P_D. The striking fact is that this embedding of P is the projection of an antipodal embedding in the sphere of the dodecahedral graph D (whence my notation P_D). Therefore, the double covering graph of P_D is D, and P_D is the projection of $(D, {}^*)$ where * is the well-known antipodal involution on D.

Now let's think about $-P$. It does not embed in the projective plane, because it has two vertex-disjoint negative circles. It does embed in U_4, as in FIgure 16.2. I will first show that the Klein bottle is too small for it. A walk around a face boundary must preserve orientation, so it is a positive closed walk; that is, it has even length. The shortest positive closed walk is a hexagon; therefore, the least possible length of a face boundary walk is 6. Euler's polyhedral formula implies that the number of faces of $-P$ embedded in U_2 is 5. The sum of all face boundary lengths is $2|E| = 30$. Thus, the only way to embed $-P$ in U_2 is to have five hexagonal faces, bounded by hexagons of P. This is where the difficulty arises. Two hexagons of P can only be face boundaries in the same embedding if they do not share a path of length 2; if they do, the middle vertex cannot have a third edge. A hexagon shares such a path with all but three other hexagons (see Lemma 3.3). It follows that there is no embedding of $-P$ in the Klein bottle.

To prove $-P$ cannot embed in U_3 I employ the circle compatibility graph from Section 3. Once more, every face boundary has even length. A face boundary of length < 10 has to be a circle, because of the girth of P. By Euler's formula there are $f = 4$ faces. If $f_l =$ the number of faces with boundary length l, $6f + 2(f_8 + 2f_{10} + 3f_{12} + 4f_{14} + \cdots) = 30$. We deduce that $f_8 + 2f_{10} + 3f_{12} + 4f_{14} + \cdots = 3$. This rules out faces of degree 14 or more, and gives the following possibilities for the face vector $(f_6, f_8, f_{10}, f_{12})$:

$$(3, 0, 0, 1), \quad (2, 1, 1, 0), \quad (1, 3, 0, 0).$$

The compatibility graph excludes all, because every face vector requires at least a triple of compatible hexagons and octagons, but there are no triangles in their compatibility graph. It follows that $-P$ has no orientation embedding in U_3.

One other Petersen signed graph, which nicely illustrates a different argument for demigenus, is $P(e)$, with one negative edge, say e. Since $P \setminus e$ is not planar, but $P(e) \setminus e$ is balanced, that part of $P(e)$ needs a torus in order to embed. Adding the edge e can be done in the torus, but to make it orientation reversing requires a crosscap. Therefore U_3 is the minimal surface for orientation embedding of $P(e)$. □

Problem 16.1. Orientation embed a signed graph like P_D in the projective plane $\mathbb{P}$. By shifting vertices around one can get a different signature, which is switching equivalent to the original, because according to the boundary-crossing rule for edge signs, pushing a vertex across the boundary switches that vertex. Is it possible to get every switching of Σ by moving the vertices around in $\mathbb{P}$? A different statement is: Given an orientation embedding $\Sigma \hookrightarrow \mathbb{P}$ and a switching function τ, is it possible to find a noncontractible, non-self-intersecting curve γ in $\mathbb{P}$ such that the edges crossing γ are precisely the negative ones? This is unknown.

I would like to know the answer, and also how to generalize the question to higher unorientable surfaces.

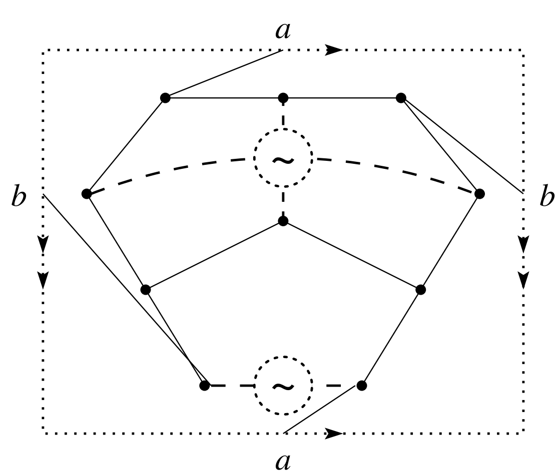

Figure 16.2. $-P$ embedded in U_4. The boundary polygon of the torus is outside and there are two crosscaps inside.

17. Johnson Quotients

In Section 15 we saw that a switching class of signed graphs is the quotient of an ordinary graph $\widetilde{\Gamma}$ modulo a fixed-point-free involutory automorphism *. I will now show that K_P is an interesting quotient graph.

The *Johnson graph* $J(m,l)$ is the graph with vertex set $\mathcal{P}_l([m])$, in which vertices A, B are adjacent when they differ in only one element, i.e., when $|A \cap B| = l-1$ ([15, p. 300], [12, Sect. 1.6]). Identifying a triple in $V(K_m) = [m]$ with the triangle it supports, $J(m,3)$ is the *triangle graph* of K_m: the graph of triangles in K_m, in which two triangles are adjacent if they have a common edge. Let's look at two of these graphs $J(m,3)$.

Consider $J(5,3)$ first. Two triples, $\{u,v,w\}$ and $\{x,y,z\}$, are adjacent if and only if their complementary vertex sets, which are edges of K_5, have empty intersection. Thus by a natural isomorphism $J(5,3) \cong \overline{L(K_5)} = P$.

Now comes the surprising part. In $J(6,3)$ write a triple $A = \{6,i,j\}$ containing 6 in terms of its complement $T := [6] \setminus A$, as $-T$, and write T as $+T$. T is a subset of [5]; so the vertices of $J(6,3)$ are the signed triples $+T$ and $-T$ for $T \in V(J(5,3))$. It is easy to check that, when $T, T' \in V(J(5,3))$, then in $J(6,3)$ either $+T, +T'$ and $-T, -T'$ are two pairs of adjacent triples (if $|T \cap T'| = 2$), or else $+T, -T'$ and $-T, +T'$ are two pairs of adjacent triples (if $|T \cap T'| = 1$). Thus, $J(6,3)$ expressed this way is the signed covering graph of a signed graph; we call the signed graph T_3.

For any two triples T, T', one of these two situations holds. Thus we are signing the complete graph K_{10} on vertex set $\mathcal{P}_3([5])$ using the following rule: An edge TT' is positive if $|T \setminus T'| = 1$, and negative if $|T \setminus T'| = 2$. The resulting signed graph T_3 is the quotient of $J(6,3)$ modulo its antipodal automorphism $T \leftrightarrow [6] \setminus T$.

Now compare the graph $J(6,3)$ with the signed graph T_3 on $\mathcal{P}_3([5])$. It is clear that $J(6,3)$ is the signed covering graph of T_3, and that $p(+T) = p(-T) = T \in V(J(5,3))$. Furthermore, the positive part of T_3 is $J(5,3)$, or equivalently $L(K_5)$, which is $\bar{P}$. Thus we have the promised interesting interpretation of K_P.

Proposition 17.1. *The signed graph K_P is (naturally isomorphic to) T_3 and its signed covering graph $\widetilde{K}_P$ is (naturally isomorphic to) $J(6,3)$.* □

Here is the natural generalization of the construction of $\mathrm{T}_3 = K_P$ from $J(6,3)$. In $J(2m,m)$ there is a fixed-point-free, involutory automorphism by complementing m-element sets. Therefore, $J(2m,m)$ is a double covering graph of a signed graph T_m. The underlying graph $|\mathrm{T}_m|$ is sometimes called the *even graph* [15, Section 9.8]. It can be defined as having vertex set $\mathcal{P}_m([2m-1])$ and edges AB if $|A \setminus B| = 1$ or $|A \cap B| = 1$. Define $J(2m-1,m,1)$ to have the same vertices as $J(2m-1,m)$, with vertices adjacent when their intersection is a singleton. Then $|\mathrm{T}_m| = J(2m-1,m) \cup J(2m-1,m,1)$ (which is K_{10} if $m=3$ but not otherwise). But what is the signature? The argument for T_3 shows that $\sigma(AB) := +1$ if $|A \setminus B| = 1$, and -1 if $|A \cap B| = 1$. Thus,

Proposition 17.2. *For $m \geq 2$, the Johnson graph $J(2m,m)$ is the signed covering graph of the signed graph T_m defined by*

$$\mathrm{T}_m^+ = J(2m-1,m) \quad \textit{and} \quad \mathrm{T}_m^- = J(2m-1,m,1).$$

□

The cases $m = 1, 2$ are exceptional in not being simple graphs. Since $J(3,2) \cong J(3,2,1) \cong K_3$, $\mathrm{T}_2 \cong (+K_3) \cup (-K_3)$. Because $J(2,1) \cong K_2$, the quotient graph $\mathrm{T}_1 \cong J(2,1)/^*$ has one

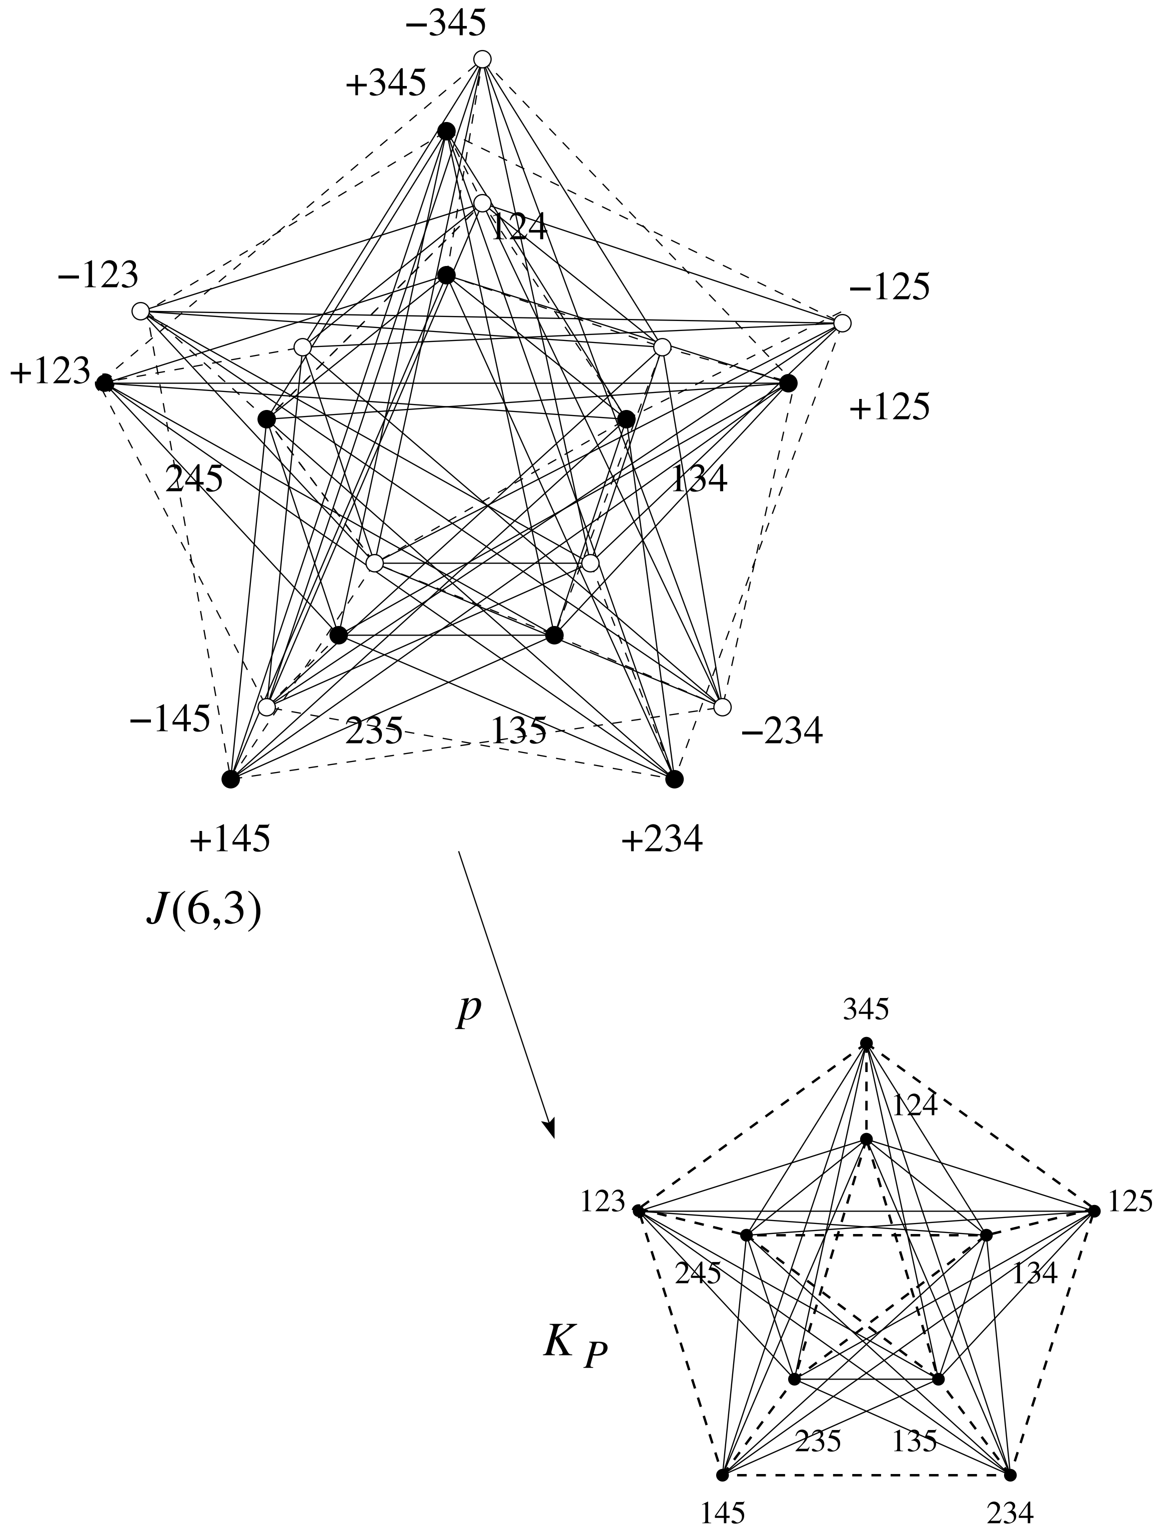

Figure 17.1. The construction of T_3, the signed K_{10}, from $J(6,3)$. Between ijk, $i'j'k' \in \mathcal{P}_3([5])$, there is a positive edge when $|\{i,j,k\} \cap \{i',j',k'\}| = 2$ and a negative edge when $|\{i,j,k\} \cap (-\{i',j',k'\})| = 2$.

vertex and one half edge. (I apologetically refer the reader to [30] for an explanation of half edges, which would lead us far astray. Suffice it to say that a half edge appears in Σ when a single edge exists between $\widetilde{v}$ and $\widetilde{v}^*$ in $\widetilde{\Sigma}$.)

PART VI. THE END

18. THE FUTURE

There have been attempts to generalize the Petersen graph, but as a rule they lack the properties that make P interesting. Working on signed graphs inspired by P led me to believe that the simplest generalization is the best. Thus, the next step should be to investigate the complement of the line graph of K_p, i.e., $\overline{L(K_p)} = J(n, 2, 0)$, or $J(2k+1, k, 0)$ (Section 17). The obvious signed graphs, the all-negative ones and the signed complete graphs, should be joined by some generalization, not evident to me, of P_D.

The strongest reasons I think $\overline{L(K_p)}$ and $J(2k+1, k, 0$ are good directions to go in are switching automorphisms and eigenvalue properties. Both seem ripe for generalization to $p > 5$. Other questions like frustration index will be harder, but tempting.

19. THE PAST

Signed graphs were invented by Frank Harary in [13] in order to handle both positive and negative relations in a small group of people, a question he studied along with the psychologist Dorwin Cartwright [6]. Balance was postulated to be a stable state, with an unbalanced state tending towards balance. it is debatable how effective this model has been; there is an extensive literature and new ideas are still appearing, some of them quite interesting mathematically.

An oddity of mathematical history is that Denes König came quite close to defining signed graphs in his pioneering textbook [16]. He had the notion of a distinguished edge set (the negative edge set) and distinguishing between circles (his word: *Kreis*) which meet that set an even or an odd number of times. He also had switching in the form of set summation with a cut (since switching X is equivalent to replacing E^- by the set sum (the symmetric difference) $E^- \oplus E(X, X^c)$). He even proved Harary's fundamental theorem of balance (see Section 5). But he never took the fundamental conceptual step of assigning multiplicative signs to edges. That is more essential than it may seem at first, and this is why I say he did not invent signed graphs.

The frustration index was introduced by the social psychologists Abelson and Rosenberg [1]. The l in the notation is from Harary's original name for it, "line index of balance".

About two decades later, the physicist Gérard Toulouse rediscovered signed graphs in connection with the general Ising model in the theory of spin glasses [26]. He modeled a spin glass as a signed graph Σ together with a changeable "state" $s : V \to \{+1, -1\}$. With respect to the state, an edge is colorfully called *satisfied* if $\sigma(e_{ij}) = s(v_i)s(v_j)$ and *frustrated* otherwise. The number of frustrated edges corresponds to the energy of the state; a *ground state* has the fewest such edges. In terms of switching, the number of frustrated edges in state s equals the number of negative edges in the switched signed graph σ^s. For more, see Section 6.

The first person to carry out switching in a way that clearly connects with signed graphs (although he did not make that connection) was J.J. Seidel in numerous papers [22]. Seidel developed the technique that is now known as Seidel switching; see Section 4. He simultaneously introduced the matrix now known as the Seidel matrix of a graph (that is, the adjacency matrix of K_Γ; see Section 13) and showed how valuable it is in the study of strongly regular graphs and regular two-graphs; see [21] and Section **??**. Seidel was also one of the authors of the fundamental paper on line graphs and their eigenvalue bound [5], in which signed graphs implicitly play an essential role (see Section **??**).

Explicit treatment of signed graphs in the context of eigenvalue bounds, however, had to await the work of G.R. Vijayakumar and his collaborators in [27, 7, 28] et al.

20. Acknowledgements

The idea of writing about the Petersen graph, and especially Section 17, was inspired in part by conversations with Jay Bagga. I thank Chris Godsil for a gentle reminder in re Petersen two-graphs and Patrick Doreian [10] for computing some quite valuable clusterability numbers, using the mathematical social psychology package Pajek [17].

References in the form Name (datex) *are to entries in* [33].

References

[1] Robert P. Abelson and Milton J. Rosenberg, Symbolic psycho-logic: a model of attitudinal cognition. *Behavioral Sci.* **3** (1958), 1–13.

[2] Norman Biggs, *Algebraic Graph Theory.* Cambridge Math. Tracts, No. 67. Cambridge University Press, London, 1974. MR 50 #151. Zbl. 284.05101. Second edn. (Exercise 19A has become Additional Result 19a), Cambridge Math. Library. Cambridge University Press, Cambridge, 1993. MR 95h:05105. Zbl. 797.05032.

[3] Andries E. Brouwer, Graph descriptions: The Desargues graph. `http://www.win.tue.nl/ãeb/graphs/Desargues.html`

[4] A.E. Brouwer, A.M. Cohen, and A. Neumaier, *Distance-Regular Graphs.* Ergebnisse der Mathematik, Third Ser., Vol. 18. Springer-Verlag, Berlin, 1989. MR 90e:05001. Zbl. 747.05073.

[5] P.J. Cameron, J.M. Goethals, J.J. Seidel, and E.E. Shult, Line graphs, root systems, and elliptic geometry. *J. Algebra* **43** (1976), 305–327. MR 56 #182. Zbl. 337.05142. Reprinted in [22, pp. 208–230].

[6] Dorwin Cartwright and Frank Harary, Structural balance: a generalization of Heider's theory. *Psychological Rev.* **63** (1956), 277–293. Reprinted in: Dorwin Cartwright and Alvin Zander, eds., *Group Dynamics: Research and Theory*, second ed., pp. 705–726. Harper and Row, New York, 1960. Also reprinted in: Samuel Leinhardt, ed., *Social Networks: A Developing Paradigm*, pp. 9–25. Academic Press, New York, 1977.

[7] P.D. Chawathe and G.R. Vijayakumar, A characterization of signed graphs represented by root system D_∞. *European J. Combin.* 11 (1990), 523–533. MR 91k:05071. Zbl. 764.05090.

[8] CombinaTexas: Combinatorics in the South-Central U.S., `http://http://www.math.tamu.edu/∼cyan/combinatexas/`

[9] James A. Davis, Structural balance, mechanical solidarity, and interpersonal relations. *Amer. J. Sociology* 68 (1963), 444–463. Reprinted with minor changes in: Joseph Berger, Morris Zelditch, Jr., and Bo Anderson, eds., *Sociological Theories in Progress, Vol. One*, Ch. 4, pp. 74–101. Houghton Mifflin, Boston, 1966. Also reprinted in: Samuel Leinhardt, ed., *Social Networks: A Developing Paradigm*, pp. 199–217.

[10] Patrick Doreian, personal communication, 1 June 2009.

[11] Patrick Doreian, Vladimir Batagelj, and Anuška Ferligoj, *Generalized Blockmodeling.* Structural Analysis in the Social Sciences, No. 25. Cambridge Univ. Press, Cambridge, Eng., 2005.

[12] Chris Godsil and Gordon Royle, *Algebraic Graph Theory.* Grad. Texts in Math., Vol. 207. Springer, New York, 2001. MR 2002f:05002. Zbl. 968.05002.

[13] F. Harary, On the notion of balance of a signed graph. *Michigan Math. J.* **2** (1953–54), 143–146 and addendum preceding p. 1. MR 16, 733h. Zbl. 56, 421c (e: 056.42103).

[14] ——, On local balance and N-balance in signed graphs. *Michigan Math. J.* **3** (1955–56), 37–41. MR 17, 394c. Zbl. 70, 185b (e: 070.18502).

[15] D.A. Holton and J. Sheehan, *The Petersen Graph.* Australian Mathematical Society Lecture Series, 7. Cambridge University Press, Cambridge, 1993. MR 94j:05037. Zbl. .

[16] Dénes König, *Theorie der endlichen und unendlichen Graphen.* Mathematik und ihre Anwendungen, Band 16. Akademische Verlagsgesellschaft, Leipzig, 1936. Reprinted by Chelsea, New York, 1950. MR 12, 195. Zbl. 13, 228 (e: 013.22803). English trans. by Ronald C. Read, *Theory of Finite and Infinite Graphs.* Birkhäuser, Boston, 1990. MR 91f:01026. Zbl. 695.05015.

[17] Wouter de Nooy, Andrej Mrvar, and Vladimir Batagelj, *Exploratory Social Network Analysis with Pajek.* Structural Anal. Soc. Sci., No. 27. Cambridge Univ. Press, Cambridge, Eng., 2005.

[18] Gerhard Ringel, The combinatorial map color theorem, *J. Graph Theory* **1** (1977), 141–155. MR 56 #2860. Zbl. 386.05030.

[19] Bruce E. Sagan, *The symmetric group: Representations, combinatorial algorithms, and symmetric functions.* Wadsworth & Brooks/Cole Math. Ser. Wadsworth & Brooks/Cole, Pacific Grove, Cal., 1991. Zbl. 823.05061.

[20] E. Sampathkumar, Point signed and line signed graphs. *Nat. Acad. Sci. Letters (India)* 7 (1984), 91–93. Zbl. 552.05051.

[21] J.J. Seidel, A survey of two-graphs. In: *Colloquio Internazionale sulle Teorie Combinatorie* (Roma, 1973), Tomo I, pp. 481–511. Atti dei Convegni Lincei, No. 17. Accad. Naz. Lincei, Rome, 1976. MR 58 #27659. Zbl. 352.05016. Reprinted in [22, pp. 146–176].

[22] ——, *Geometry and Combinatorics: Selected Works of J.J. Seidel.* D.G. Corneil and R. Mathon, eds. Academic Press, Boston, 1991. MR 92m:01098. Zbl. 770.05001.

[23] Daniel Slilaty, Projective-planar signed graphs and tangled signed graphs. *J. Combin. Theory Ser. B* 97 (2007), no. 5, 693–717. MR 2008j:05161. Zbl. 1123.05046.

[24] Tadeusz Sozański, Processus d'équilibration et sous-graphes équilibrés d'un graphe signé complet. *Math. Sci. Humaines*, No. 55 (1976), 25–36, 83. MR 58 #27613.

[25] S. Stahl, Generalized embedding schemes. *J. Graph Theory* **2** (1978), 41–52. MR 58 #5318. Zbl. 396.05013.

[26] Gérard Toulouse. Theory of the frustration effect in spin glasses: I. *Commun. Phys.* 2 (1977), 115–119. Reprinted in: M. Mézard, G. Parisi, and M.A. Virasoro, *Spin Glass Theory and Beyond*, pp. 99–103. World Scientific Lecture Notes in Physics, Vol. 9. World Scientific, Singapore, 1987.

[27] G.R. Vijayakumar, Signed graphs represented by D_∞. *European J. Combin.* 8 (1987), 103–112. MR 88b:05111. Zbl. 678.05058.

[28] ——, Representation of signed graphs by root system E_8. *Graphs Combin.* 10 (1994), 383–388. MR 96a:05128. Zbl. 821.05040.

[29] Thomas Zaslavsky, Characterizations of signed graphs. *J. Graph Theory* **5** (1981), 401–406. MR 83a:05122. Zbl. 471.05035.

[30] ——, Signed graphs. *Discrete Appl. Math.* **4** (1982), 47–74. Erratum. *Discrete Appl. Math.* **5** (1983), 248. MR 84e:05095. Zbl. 503.05060.

[31] ——, Orientation embedding of signed graphs. *J. Graph Theory* **16** (1992) 399–422. MR 93i:05056. Zbl. 778.05033.

[32] ——, Matrices in the theory of signed simple graphs. In: B.D. Acharya, G.O.H. Katona, and J. Nesetril, eds., *Advances in Discrete Mathematics and Its Applications* (Proc. Int. Conf. Discrete Math., Mysore, India, 2008). Ramanujan Math. Soc. Lecture Notes Ser. Math. Submitted.

[33] ——, A mathematical bibliography of signed and gain graphs and allied areas. *Electronic J. Combin.*, Dynamic Surveys in Combinatorics (1998), No. DS8 (electronic). MR 2000m:05001a. Zbl. 898.05001. Current draft edition on the WorldWideWeb at
`http://www.math.binghamton.edu/zaslav/Bsg/`

DEP'T OF MATHEMATICAL SCIENCES, BINGHAMTON UNIVERSITY OF SUNY, BINGHAMTON, NY 13902-6000, U.S.A.